arXiv:1409.6326v5 [math.GR] 5 Sep 2016

CHARACTERISATIONS OF ALGEBRAIC PROPERTIES OF GROUPS IN TERMS OF HARMONIC FUNCTIONS

MATTHEW C. H. TOINTON

Abstract. We prove various results connecting structural or algebraic properties of graphs and groups to conditions on their spaces of harmonic functions. In particular: we show that a group with a finitely supported symmetric measure has a finite-dimensional space of harmonic functions if and only if it is virtually cyclic; we present a new proof of a result of V. Trofimov that an infinite vertex-transitive graph admits a non-constant harmonic function; we give a new proof of a result of T. Ceccherini-Silberstein, M. Coornaert and J. Dodziuk that the Laplacian on an infinite, connected, locally finite graph is surjective; and we show that the positive harmonic functions on a non-virtually nilpotent linear group span an infinite-dimensional space.

Contents

1. Introduction 1
2. Background and notation 6
3. A Garden of Eden theorem 10
4. Transpose-harmonic functions and surjectivity of Laplacians 11
5. A duality result for harmonic functions 12
6. Existence of non-constant harmonic functions on graphs 13
7. Harmonic functions on virtually abelian groups 17
8. Positive harmonic functions on linear groups 18
9. Random walks on virtually cyclic groups 20
10. Harmonic functions on groups with virtually cyclic quotients 21
11. Groups with finite-dimensional spaces of harmonic functions 28
Appendix A. Further applications of our Garden of Eden theorem 29
References 31

1. Introduction

One can often obtain algebraic information about a group by considering it as a geometric object. For example, if G is a group and $S \subset G$ is a finite, symmetric set then one can construct the *Cayley graph* (G, S) of G with respect to S by declaring the elements of G to be vertices and saying that

2010 *Mathematics Subject Classification.* 20F65 (primary), 60G50 (secondary).

Key words and phrases. Discrete harmonic function, discrete Laplacian, random walk, Cayley graph, linear cellular automaton.

The majority of this work was supported by a Junior Research Fellowship from Homerton College, University of Cambridge. This work was also partially supported by a grant from the Cambridge Philosophical Society.

x and y are joined by an edge if and only if there is some non-identity element $s \in S$ such that $xs = y$.

One way of studying the geometry of a Cayley graph, or indeed any graph, is to consider the behaviour of probabilistic processes on it. In this paper we are particularly concerned with linking the algebra and geometry of groups and graphs to spaces of *harmonic functions* on them.

Before we define these, let us establish some notation. A *weighted graph* Γ is a graph in which to each edge xy we associate a real number $\omega_{xy} = \omega_{yx} > 0$ called a *weight*; the degree of a vertex x is then given by $\deg x = \sum_{y \sim x} \omega_{xy}$. We define the *Laplacian* $\Delta = \Delta_\Gamma$ on Γ by setting $\Delta f(x) = f(x) - \frac{1}{\deg x} \sum_{y \sim x} \omega_{xy} f(y)$ for every function $f : \Gamma \to \mathbb{R}$.

If G is a group then a probability measure μ on G is said to be a *generating probability measure* if the semigroup generated by its support, $\operatorname{supp} \mu$, is G; it is said to be *symmetric* if $\mu(g) = \mu(g^{-1})$ for every $g \in G$. For a group G with a finitely supported generating probability measure μ we write $\Delta = \Delta_\mu$ for the *Laplacian* on G with respect to μ, defined by setting $\Delta f(x) = f(x) - \sum_{s \in \operatorname{supp} \mu} \mu(s) f(xs)$ for every function $f : G \to \mathbb{R}$. Note that if μ is symmetric and Γ is the Cayley graph of G with respect to $\operatorname{supp} \mu$, weighted such that $\omega_{xy} = \mu(x^{-1}y)$, then $\Delta_\mu = \Delta_\Gamma$. We denote this weighted Cayley graph by (G, μ).

A *harmonic* function on a weighted graph Γ or group G with generating probability measure μ is defined to be a function belonging to the kernel of the corresponding Laplacian. We write $H(G, \mu)$ for the space of harmonic functions on G with respect to μ.

Perhaps the most famous example of a result linking the algebraic structure of a group to the geometry of a Cayley graph is M. Gromov's celebrated theorem on groups of polynomial growth, which states that a certain *geometric* condition on a Cayley graph (G, S) (polynomial volume growth) is characteristic of a certain *algebraic* condition on the subgroup of G generated by S (virtual nilpotency) [13]. A recent proof of Gromov's theorem due to B. Kleiner [17] provides an example of how harmonic functions are related to the algebra and geometry of groups, since a key step in Kleiner's proof is to show that if a group has polynomial growth then the vector space of harmonic functions on (G, S) that grow at most linearly in the Cayley-graph distance from the identity is finite dimensional.

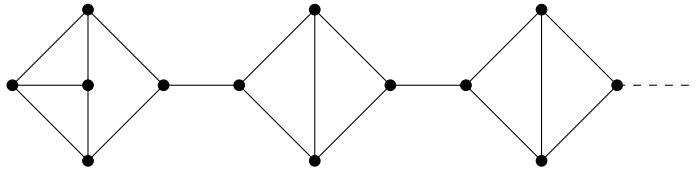
While Kleiner's proof of Gromov's theorem essentially uses the space of linearly growing harmonic functions as a tool to characterise an algebraic condition on a group in terms of a geometric condition, in principle it should be possible to characterise certain algebraic or geometric conditions purely in terms of spaces of harmonic functions. Indeed, in a very recent preprint, T. Meyerovitch and A. Yadin [19] have shown that in the case of a finitely generated group that is linear or virtually soluble, being virtually nilpotent is *equivalent* to having a finite-dimensional space of linearly growing harmonic functions. This equivalence is, moreover, conjectured to hold for all finitely generated groups [19].

The first result of the present paper shows that finite-dimensionality of the space of *all* harmonic functions on a group is also equivalent to a simple algebraic condition.

Theorem 1.1. *Let G be an infinite group, and let μ be a symmetric, finitely supported generating probability measure on G. Then the space of harmonic functions on (G, μ) is finite dimensional if and only if G contains a finite-index subgroup isomorphic to $(\mathbb{Z}, +)$.*

One can also consider, as Kleiner did in his proof of Gromov's theorem, subspaces of $H(G, \mu)$ consisting of functions of polynomial growth. Given a group G with a finitely supported generating probability measure μ, denote by $|g| = |g|_\mu$ the word distance of g from the identity with respect to the generating set $\operatorname{supp} \mu$. The space $H^k(G, \mu)$ of harmonic functions on G of polynomial growth of

FIGURE 1. An infinite regular graph with no non-constant harmonic functions [24, Remark 2]



degree at most k is then defined by $H^k(G,\mu) = \{h \in H(G,\mu) : |h(x)| \ll_h |x|^k \text{ as } x \to \infty\}$. The set $\bigcup_{k=1}^{\infty} H^k(G,\mu)$ of all harmonic functions of polynomial growth on G is of course also a subspace of $H(G,\mu)$. Emmanuel Breuillard has pointed out a result of Osin [21] that combines with the proof of Theorem 1.1 and the result of Meyerovitch and Yadin [19] to give the following stronger statement, valid for all groups except perhaps those that are amenable but not elementary amenable.

Corollary 1.2. *Let G be an infinite group that is either elementary amenable or non-amenable, and let μ be a symmetric, finitely supported generating probability measure on G. Then the space of all harmonic functions of polynomial growth on (G,μ) is finite dimensional if and only if G contains a finite-index subgroup isomorphic to $(\mathbb{Z},+)$.*

Conjecture 1.3. *Corollary 1.2 holds for all finitely generated groups.*

A fairly immediate consequence of Theorem 1.1 (at least in the presence of other, standard, results) is that the space of harmonic functions on a group with a symmetric, finitely supported generating probability measure is 1-dimensional if and only if the group is finite. In fact, this was already well known, and there exist far simpler proofs than via Theorem 1.1. Indeed, V. Trofimov [24] shows that this characterisation holds, more generally, for vertex-transitive graphs. The second result of this paper is a new proof of Trofimov's result, valid in the even more general setting of vertex-transitive weighted graphs (Trofimov's proof could conceivably also work in this more general setting).

Proposition 1.4. *Let Γ be an infinite, locally finite, vertex-transitive weighted graph. Then Γ admits a non-constant harmonic function.*

In Section 6 we prove Proposition 1.4 in the case that the random walk on Γ is transient; for the recurrent case we refer the reader to [24] (where the two cases are also treated separately). See Section 2 for definitions of *transient*, *recurrent* and *random walk*.

Remarks 1.5. Trofimov's result is in fact stronger than Proposition 1.4, as it proves the existence of a function whose growth rate is bounded in terms of the rate of volume growth of metric balls in Γ. Nonetheless, it seems to be of interest to have an alternative proof of the qualitative statement, and in any case we deduce Proposition 1.4 fairly immediately from a slightly more general result (Proposition 6.1, below) that is an important ingredient in our proof of Theorem 1.1.

Proposition 1.4 does not necessarily hold if Γ is not vertex transitive, as can be seen by considering the graph in Figure 1. This example was presented explicitly in a talk of Coornaert (available at `http://www-irma.u-strasbg.fr/~coornaer/florence-laplacian-2012.pdf`), having been observed by Trofimov [24, Remark 2].

The results we have stated so far characterise algebraic conditions on a group in terms of the kernel of the Laplacian. Our next result characterises a structural condition on a graph in terms of the *image* of the Laplacian.

A simple rank-nullity argument shows that the Laplacian on a finite graph is not surjective, since its kernel contains the constant functions. T. Ceccherini-Silberstein, M. Coornaert and J. Dodziuk [9, Theorem 1.1] show that the converse is also true for connected graphs. In Section 4 we give a new proof of this result, valid in weighted graphs (Ceccherini-Silberstein, Coornaert and Dodziuk's proof could conceivably also work in this more general setting).

Proposition 1.6. *The Laplacian on an infinite, connected, locally finite weighted graph Γ is surjective onto $\mathbb{R}^\Gamma$. Thus, the Laplacian on a locally finite weighted graph Γ is surjective onto $\mathbb{R}^\Gamma$ if and only if every connected component of Γ is infinite.*

Our proof of Proposition 1.6 is inspired by an earlier, less general, result of Ceccherini-Silberstein and Coornaert, which states that the Laplacian on an infinite Cayley graph is surjective [6, Theorem 1.1]. We also show that this holds for more general Laplacians on groups.

Proposition 1.7. *Let G be an infinite group, and let μ be a (not necessarily symmetric) finitely supported generating probability measure on G. Then Δ_μ is surjective.*

It is quite likely that the argument of Ceccherini-Silberstein and Coornaert [6] could also give a proof of Proposition 1.7. However, our proof of Proposition 1.7 is simpler than the argument of [6] (see Remarks 4.3, below), and since Proposition 1.7 cannot be concluded directly from either Proposition 1.6 or [6, Theorem 1.1], it seems, in any case, worth recording a proof here.

An important tool in this paper is the so-called *Garden of Eden theorem* for linear cellular automata, originally due to Ceccherini-Silberstein and Coornaert [5]. Given its importance to our arguments, we introduce it briefly here.

If G is a group and A is a set, called the *alphabet*, then G acts on the set A^G of maps $f : G \to A$ via $g \cdot f(x) = f(g^{-1}x)$. If $f : G \to A$ and $M \subset G$ then we denote by $f|_M$ the restriction of f to M. A *cellular automaton* over G on the alphabet A is a map $\tau : A^G \to A^G$ with the property that there is some finite set $M \subset G$ and a map $\lambda : A^M \to A$ such that $\tau(f)(g) = \lambda((g \cdot f)|_M)$. The set M is called a *memory set* for τ, and λ is called a *local defining map*.

Given an initial state $f_0 \in A^G$, one can consider τ as defining a dynamical process on A^G by setting $f_{i+1} = \tau(f_i)$ to obtain a sequence $f_0, f_1, f_2, \ldots$ of configurations in A^G. A configuration $f \in A^G$ is then said to be a *Garden of Eden* configuration if it is not in the image of τ, and hence can appear only as an initial configuration in this dynamical process.

The term *Garden of Eden theorem* for a class of cellular automata is often used to describe a result giving a necessary and sufficient condition for the existence of Garden of Eden configurations, or, to put it another way, a necessary and sufficient condition for a cellular automaton in the class to be surjective. There are various results depending on the alphabet and the group; we refer the reader to [5, 8] for more detailed background to this area.

The class of interest to us is the class of *linear cellular automata*, in which A is a finite-dimensional vector space $V = \mathbb{K}^r$ over a field $\mathbb{K}$ and a linear cellular automaton is a cellular automaton that is also a linear map $V^G \to V^G$.

Theorem 1.8 (Garden of Eden theorem for linear cellular automata; Ceccherini-Silberstein–Coornaert [8, Theorem 8.9.6])**.** *Let V be a finite-dimensional vector space and let G be an amenable group. Then a linear cellular automaton $\tau : V^G \to V^G$ is surjective if and only if it is pre-injective.*

Here, and throughout this paper, a linear map on V^G is said to be *pre-injective* if its restriction to the subspace V_0^G of finitely supported functions in V^G is injective.

The Laplacian on a group with a finitely supported generating probability measure is an example of a linear cellular automaton, and so Theorem 1.8 can readily be applied to such a Laplacian, provided that the group is amenable. However, in this paper we are concerned with Laplacians on arbitrary groups, and even graphs, and so we seek a version of Theorem 1.8 that holds in this greater generality.

Given a locally finite graph Γ and an alphabet A, we say that a map $\tau : A^\Gamma \to A^\Gamma$ is *locally specifiable* if $\tau(f)(x)$ depends only on $f(x)$ and $f(y)$ for $y \sim x$. Note, in particular, that if τ is a cellular automaton on a group G with memory set M then τ is a locally specifiable map on the Cayley graph $(G, M \cup M^{-1})$. The Laplacian on a locally finite graph is also locally specifiable.

The result underpinning much of this paper is the following.

Theorem 1.9. *Let V be a finite-dimensional vector space and let Γ be a locally finite graph. Then a locally specifiable linear map $\tau : V^\Gamma \to V^\Gamma$ is surjective if and only if its transpose τ' is pre-injective.*

Here the *transpose* of τ is defined in terms of the natural (possibly infinite) matrix representation of τ, which we define precisely in Section 2. The transpose of τ is then simply the locally specifiable linear map whose corresponding matrix is the transpose of the matrix corresponding to τ.

Let us emphasise here that Theorem 1.9 applies, in particular, to linear cellular automata over non-amenable groups, and is, in that sense, considerably more general than Theorem 1.8.

Corollary 1.10 (Garden of Eden theorem for symmetric linear cellular automata over non-amenable groups)**.** *Let V be a finite-dimensional vector space and let G be a (not necessarily amenable) group. Then a symmetric linear cellular automaton $\tau : V^G \to V^G$ is surjective if and only if it is pre-injective.*

Remarks 1.11. The reader may refer to [5, §5] or [8, §8.10-8.11] for examples of (asymmetric) linear cellular automata on non-amenable groups for which Theorem 1.8 fails; thus, generalisations to non-amenable groups in the spirit of Corollary 1.10 must necessarily have some additional hypothesis on the map τ.

With a bit more work, one can adapt some of the techniques from [5] to recover Theorem 1.8 in full from Theorem 1.9; see Appendix A.

Outline of the paper. In Section 2 we give detailed definitions and present some necessary background material. In Section 3 we prove Theorem 1.9, before applying it in Section 4 to prove Propositions 1.6 and 1.7. In Section 5 we use all of these results to develop a tool for proving the existence of harmonic functions on a graph or group, and then in Section 6 we use this tool to prove a slight technical generalisation of Proposition 1.4. In Section 7 we prove the easier 'direct' direction of Theorem 1.1, and in Section 8 we prove the 'inverse' direction in the case of a non-virtually nilpotent linear group. In Sections 9–11 we reduce Theorem 1.1 to the linear case and complete the proof, as well as proving Corollary 1.2.

In the appendix we present two additional applications of Theorem 1.9. In the first, we recover Theorem 1.8; in the second, we reformulate a conjecture of I. Kaplansky, the so-called 'stable-finiteness' conjecture.

Acknowledgements. The author is grateful to Emmanuel Breuillard, Sara Brofferio, Michel Coornaert, Ben Green and Gady Kozma for helpful conversations. Thanks are also due to one anonymous referee for noticing an error in an earlier claimed proof of Theorem 1.1, and another anonymous

referee for helpful comments on an earlier version of this paper. Some of this work was carried out at the Institut Henri Poincaré, Paris, during its excellent trimester 'Random walks and asymptotic geometry of groups', 2014.

2. Background and notation

In this section we set much of our notation and present general background material from the literature. Much of this is standard; see, for example, [22].

Throughout this paper, by a *graph* Γ we mean an undirected weighted graph with no loops and no multiple edges. We denote by e some distinguished vertex; this vertex is always the identity in the case that Γ is a weighted Cayley graph. We write $x \sim y$ to indicate that x and y are neighbours. An isomorphism of weighted graphs is an isomorphism of graphs that preserves weights. A weighted graph is called *regular* if $\deg x = \sum_{y\sim x} \omega_{xy}$ is independent of the vertex x.

Denote by $d = d_\Gamma$ the graph metric on a graph Γ; thus, for vertices $x \neq y \in \Gamma$ the quantity $d(x, y)$ is equal to length of a path of minimum length joining x to y. If G is a group with a finitely supported generating probability measure μ then we denote by $d = d_\mu$ the graph distance on (G, μ) (thus d_μ is the word metric with respect to the generating set $\operatorname{supp}\mu$).

If Γ is a graph or group and $V = \mathbb{K}^n$ is a vector space then for each vertex or element $x \in \Gamma$ and each $i = 1, \dots, n$ we denote by $\delta^i_x : \Gamma \to V$ the map defined by $\delta^i_x(x) = e_i$, and $\delta^i_x(y) = 0$ for every $y \neq x$. In the event that $n = 1$ we drop the superscript and define $\delta_x : \Gamma \to \mathbb{K}$ by $\delta_x(x) = 1$ and $\delta_x(y) = 0$ for every $y \neq x$. The δ^i_x form a basis for the space V_0^Γ of finitely supported V-valued functions on Γ (and, for the purposes of this paper, should 'morally' be thought of as a basis for V^Γ).

This space V_0^Γ is invariant under any locally specifiable linear map $\tau : V^\Gamma \to V^\Gamma$, and so we may consider the (possibly infinite) matrix of the restriction $\tau|_{V_0^\Gamma}$ with respect to this basis. In fact, τ is entirely determined by its restriction to V_0^Γ, and so the matrix of $\tau|_{V_0^\Gamma}$ with respect to this basis completely determines τ. Moreover, the composition of such matrices respects the composition of the corresponding linear maps. Throughout this paper, when we refer to the matrix of a locally specifiable linear map $\tau : V^\Gamma \to V^\Gamma$ we mean the matrix of its restriction $\tau|_{V_0^\Gamma}$ with respect to the basis $\{\delta^i_x\}$.

Given a finite set Y and a function $f : Y \to \mathbb{R}$, we generally denote by $\mathbb{E}_{y\in Y}$ the average $\mathbb{E}_{y\in Y} f(y) = \frac{1}{|Y|}\sum_{y\in Y} f(y)$. However, in the specific case that μ is a generating probability measure on a group G, and S is the support of μ, given a function $f : S \to \mathbb{R}$ the notation $\mathbb{E}_{s\in S}$ means the average $\mathbb{E}_{s\in S} f(s) = \sum_{s\in S} \mu(s) f(s)$. Note that these definitions agree only if μ is the uniform probability measure on S.

If G_1, G_2 are groups and $\phi : G_1 \to G_2$ is a surjective homomorphism then given a finitely supported generating probability measure μ on G_1 we define a finitely supported generating probability measure $\phi(\mu)$ on G_2 by setting $\phi(\mu)(g) = \sum_{\overline{g}\in\phi^{-1}(g)} \mu(\overline{g})$. Note that if μ is symmetric then so is $\phi(\mu)$.

Lemma 2.1. *Let G_1 be a group with a finitely supported generating probability measure μ. Suppose that $\phi : G_1 \to G_2$ is a surjective homomorphism, and that $f : G_2 \to \mathbb{R}$. Then $f \circ \phi$ is harmonic with respect to μ if and only if f is harmonic with respect to $\phi(\mu)$.*

Proof. Given an arbitrary $g_2 \in G_2$, the surjectivity of ϕ implies that there exists $g_1 \in G_1$ such that $\phi(g_1) = g_2$. On the other hand, given an arbitrary $g_1 \in G_1$, we may simply define $g_2 \in G_2$ by

$g_2 = \phi(g_1)$. In either case, $f(g_2) = f \circ \phi(g_1)$ and $\mathbb{E}_{s\in\phi(S)} f(g_2 s) = \mathbb{E}_{s\in S} f(g_2\phi(s)) = \mathbb{E}_{s\in S} f \circ \phi(g_1 s)$, from which the lemma follows easily. □

Given a subset A of a graph Γ, or of a group G with a finitely supported generating probability measure μ, we define the *neighbourhood* A^+ of A to be the set $A^+ = \{x \in \Gamma : d(x, A) \leq 1\}$, the *interior* A° of A to be the set $A^\circ = \{x \in A : \{x\}^+ \subset A\}$, the *inner boundary* $\partial^- A$ of A to be the set $\partial^- A = A \backslash A^\circ$, and the *outer boundary* $\partial^+ A$ of A to be the set $\partial^+ A = A^+ \backslash A$.

Let Γ be a locally finite weighted graph, or a group with a finitely supported generating probability measure μ. Let A be a subset of Γ, and let D be a subset of Γ containing A^+. Then we say that a function $h : D \to \mathbb{R}$ is *harmonic on* A if we have $\Delta h(x) = 0$ for each $x \in A$.

The following is an immediate consequence of the definition of harmonicity.

Lemma 2.2 (Maximum principle)**.** *Let Γ be a locally finite graph, or a group with a finitely supported generating probability measure μ, and let A be a connected subset of Γ. Suppose that $f : A^+ \to \mathbb{R}$ is harmonic on A and achieves a maximum on A. Then f is constant.*

Harmonic functions on graphs and groups are intimately connected to random walks. Given a graph Γ and a vertex $x \in \Gamma$, the *random walk* starting at x is a sequence of Γ-valued random variables $X_0, X_1, X_2, \ldots$, with $X_0 = x$ with probability 1 and each subsequent X_n chosen from among the neighbours of X_{n-1} such that $X_n = y$ with probability $\omega_{X_{n-1}y} / \deg X_{n-1}$. Given a group G with a finitely supported generating probability measure μ, the random walk on the pair (G, μ) starting at $x \in G$ is a sequence of G-valued random variables $X_0, X_1, X_2, \ldots$, with $X_0 = x$ with probability 1 and each subsequent X_n taking the value $X_{n-1}s$ with probability $\mu(s)$. We say that the random walk on (G, μ) is *symmetric* if μ is symmetric.

Given an event B, we denote by $\mathbb{P}_x[\, B \,]$ the conditional probability $\mathbb{P}[\, B \,|\, X_0 = x \,]$. Given another event C, we denote by $\mathbb{P}_x[\, B \,|\, C \,]$ the conditional probability $\mathbb{P}[\, B \,|\, C \text{ and } \{X_0 = x\} \,]$. We use the conditional expectation notation $\mathbb{E}_x$ similarly.

If A is a subset of Γ, we write $T_A := \inf\{t : X_t \in A\}$, with $T_A = \infty$ if $X_t \notin A$ for all t. The random variable T_A is often called a *stopping time* for the random walk. If A is the singleton $\{x\}$ then we abbreviate $T_x := T_{\{x\}}$.

The next few results are standard; see, for example, [22].

Lemma 2.3 (Harmonic functions are determined by their boundary values)**.** *Let Γ be a graph, or a group with a finitely supported generating probability measure, and let A be a finite subset of Γ with non-empty outer boundary. Let $f_0 : \partial^+ A \to \mathbb{R}$. Then the function $f : A^+ \to \mathbb{R}$ defined by $f(x) = \mathbb{E}_x \left[f_0 \left(X_{T_{\partial^+ A}} \right) \right]$ is harmonic on A and agrees with f_0 on $\partial^+ A$, and is unique with respect to these two properties.*

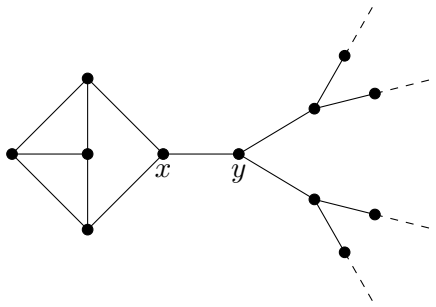
Corollary 2.4. *Let G be a group and let A be a finite subset of G. Suppose that $f_1, f_2 : A^+ \to \mathbb{R}$ are harmonic on A, and that $f_1 \geq f_2$ on $\partial^+ A$. Then $f_1 \geq f_2$ on the whole of A^+.*

Lemma 2.5. *Let x, y be vertices in a vertex-transitive weighted graph Γ. Then $\mathbb{P}_x[\, X_{2n} = y \,] \leq \mathbb{P}_e[\, X_{2n} = e \,]$.*

Remark 2.6. Lemma 2.5 does not necessarily hold if $2n$ is replaced by n. For example, if n is odd then in the Cayley graph $(\mathbb{Z}, \pm 1)$ we have $\mathbb{P}_0[\, X_n = 0 \,] = 0$.

Proposition 2.7. *Let Γ be a locally finite vertex-transitive weighted graph, and let $x, y \in \Gamma$. Then for each n we have $\mathbb{P}_x[\, T_y = n \,] = \mathbb{P}_y[\, T_x = n \,]$.*

FIGURE 2. A regular graph in which $\mathbb{P}_x[\,T_y < \infty\,] > \mathbb{P}_y[\,T_x < \infty\,]$.



Remark 2.8. Proposition 2.7 is trivial for a Cayley graph. It does not necessarily hold in a regular graph that is not vertex transitive; see Figure 2.

Proposition 2.7 seems to be well known – see, for example, [1, Proposition 2] for a proof in the case of a finite graph – but the author was unable to find in the literature a proof of it as stated, so we present one here. A key step is the following lemma.

Lemma 2.9. *Let Γ be a locally finite vertex-transitive weighted graph, and let $n \in \mathbb{N}$. Then for every $x, y \in \Gamma$ we have*

$$\begin{aligned}\mathbb{P}_x[\,X_n = x, \textit{ and } X_i \neq y \textit{ for all } i = 1, \dots, n-1\,] \\ = \mathbb{P}_y[\,X_n = y, \textit{ and } X_i \neq x \textit{ for all } i = 1, \dots, n-1\,]\end{aligned}$$

Proof. If $n = 0$ then the lemma is trivial, so by induction we may fix $n > 0$ and assume that

$$\begin{aligned}\mathbb{P}_x[\,X_r = x, &\text{and } X_i \neq y \text{ for all } i = 1, \dots, r-1\,] \\ &= \mathbb{P}_y[\,X_r = y, \text{and } X_i \neq x \text{ for all } i = 1, \dots, r-1\,] \\ &= u_r,\end{aligned}$$

say, for every $r < n$. Moreover, since Γ is regular, if $z_0, \dots, z_r$ is a path from x to y then $\mathbb{P}_x[\,X_0 = z_0, \dots, X_r = z_r\,] = \mathbb{P}_y[\,X_0 = z_r, \dots, X_r = z_0\,]$. This means, in particular, that if $v_r(x, y)$ is the probability of moving from x to y in r steps, without visiting either x or y in between, then

$$v_r(x, y) = v_r(y, x) = v_r, \tag{2.1}$$

say, for every r.

It is immediate from the vertex transitivity of Γ that we have

$$\mathbb{P}_x[\,X_n = x\,] = \mathbb{P}_y[\,X_n = y\,], \tag{2.2}$$

and so it suffices to show that we have

$$\begin{aligned}\mathbb{P}_x[\,X_n = x, \text{ and } X_i = y \text{ for some } i = 1, \dots, n-1\,] \\ = \mathbb{P}_y[\,X_n = y, \text{ and } X_i = x \text{ for some } i = 1, \dots, n-1\,].\end{aligned} \tag{2.3}$$

Given $k \geq 1$ and a sequence $0 \leq a_1 < b_1 \leq a_2 < b_2 \leq \ldots \leq a_k < b_k \leq n$ of integers, define the event $L_{x,y}(n; k; a_1, \dots, a_k; b_1, \dots, b_k)$ to be the event that $X_0 = X_n = x$ and, if $0 = t_1 < \ldots < t_l = n$ are all the times t at which $X_t \in \{x, y\}$ and we set $A = \{t_i : X_{t_i} \neq X_{t_{i+1}}\}$ and $B = \{t_{i+1} : X_{t_i} \neq X_{t_{i+1}}\}$,

then we have $A = \{a_1, \dots, a_k\}$ and $B = \{b_1, \dots, b_k\}$. Setting $a_{k+1} = n$ and $b_0 = 0$ for notational convenience, we have

$$\begin{aligned} \mathbb{P}_x[\, L_{x,y}(n; k; a_1, \dots, a_k; b_1, \dots, b_k)\,] &= \prod_{i=0}^{k} u_{a_{i+1}-b_i} \prod_{j=1}^{k} v_{b_j - a_j} \\ &= \mathbb{P}_y[\, L_{y,x}(n; k; a_1, \dots, a_k; b_1, \dots, b_k)\,]. \end{aligned} \tag{2.4}$$

However, the event $\{\, X_0 = X_n = x$, and $X_i = y$ for some $i = 1, \dots, n-1\,\}$ is precisely the disjoint union of all events $L_{x,y}(n; k; a_1, \dots, a_k; b_1, \dots, b_k)$ with $k \geq 1$, and so (2.3) follows immediately from (2.4). The lemma is then immediate from (2.2) and (2.3). □

Proof of Proposition 2.7. We prove the more precise statement that

$$\begin{aligned} &\mathbb{P}_x[\, T_y = n \text{ and } \max\{t < n : X_t = x\} = r\,] \\ &\qquad = \mathbb{P}_y[\, T_x = n \text{ and } \max\{t < n : X_t = x\} = r\,] \end{aligned}$$

for every $r \geq 0$. Indeed, this follows readily from Lemma 2.9 and (2.1), and the observation that

$$\begin{aligned} &\mathbb{P}_x[\, T_y = n \text{ and } \max\{t < n : X_t = x\} \;=\; r\,] \\ &\qquad = v_{n-r}(x, y)\mathbb{P}_x[\, X_r = x, \text{and } X_i \neq y \text{ for all } i = 1, \dots, r-1\,]. \end{aligned}$$

□

Remark 2.10. The only properties of Γ that we used in the proof of Proposition 2.7 were its regularity and (2.2). These properties are satisfied, more generally, by *walk-regular* (unweighted) graphs (see [11, 12] for definitions and background). Proposition 2.7 therefore also holds in walk-regular unweighted graphs.

A vertex x of a graph, or a group with a finitely supported generating probability measure, is called *recurrent* for the random walk on the graph or group if $\mathbb{P}_x[\, T_x < \infty\,] = 1$, and *transient* for the random walk otherwise. In the case of a connected graph or a group this is independent of the choice of vertex, and so it makes sense to define the random walk on a connected graph, or on a group with a finitely supported generating probability measure, to be *recurrent* if $\mathbb{P}_e[\, T_e < \infty\,] = 1$, and *transient* otherwise.

Write R_x for the number of times the random walk visits the vertex x. Note that in the case of a transient random walk the variable R_e has a geometric distribution under the probability measure $\mathbb{P}_e$, from which the following well-known fact easily follows.

Lemma 2.11. *The random walk on a connected graph, or on a group with a finitely supported generating probability measure, is transient if and only if* $\mathbb{E}_e[R_e] < \infty$.

In the case of a group, if we require probability measures to be symmetric then recurrence or transience of the random walk is even independent of the choice of finitely supported generating probability measure [26, Proposition 4.2]. It therefore makes sense simply to define a finitely generated group to be *recurrent* if some symmetric random walk on it is recurrent, and *transient* otherwise.

N. Varopoulos has characterised those groups that are recurrent.

Proposition 2.12 (Varopoulos [25, 26])**.** *Let G be a group with a symmetric, finitely supported generating probability measure μ. Then the random walk on (G, μ) is recurrent if and only if G is finite or has a finite-index subgroup isomorphic to $\mathbb{Z}$ or $\mathbb{Z}^2$.*

We close this section by recording the following standard but repeatedly useful reduction.

Lemma 2.13. *Let G be a group and let H be a finite-index subgroup of G. Then there exists a finite-index subgroup $H' < H$ that is normal in G.*

Proof. It is easy to verify that the subgroup $H' = \bigcap_{gH \in G/H} gHg^{-1}$ is well defined, normal and of finite index in G. □

3. A Garden of Eden theorem

In this section we prove Theorem 1.9. Throughout this section, we write e for an arbitrary distinguished vertex of the graph Γ under consideration, and write $B(n) = B_e(n)$ for the ball of radius n about e.

Lemma 3.1 (Ceccherini-Silberstein–Coornaert)**.** *Let Γ be a connected, locally finite graph and let $\tau : V^\Gamma \to V^\Gamma$ be a locally specifiable linear map. Suppose that $f : \Gamma \to V$ is such that for every n there is a function $v_n : \Gamma \to V$ such that $\tau(v_n)$ and f agree on the ball $B(n)$. Then there is a function $w : \Gamma \to V$ such that $f = \tau(w)$.*

Proof [5, Lemma 3.1]. For each $n \geq 2$, denote by τ_n the linear map $V^{B(n)} \to V^{B(n-1)}$ induced by τ, and define L_n to be the affine subspace of $V^{B(n)}$ given by $L_n = \tau_n^{-1}(f|_{B(n-1)})$. Note in particular that $v_{n-1}|_{B(n)} \in L_n$, so that L_n is non-empty.

For $n \leq m$, the restriction map $V^{B(m)} \to V^{B(n)}$ induces an affine map $\pi_{n,m} : L_m \to L_n$, and so we may define an affine subspace $K_{n,m} \subset L_n$ by $K_{n,m} = \pi_{n,m}(L_m)$. Since

$$\pi_{n_1,n_3} = \pi_{n_1,n_2} \circ \pi_{n_2,n_3} \tag{3.1}$$

whenever $n_1 \leq n_2 \leq n_3$, for any fixed n we have $K_{n,n} \supset K_{n,n+1} \supset K_{n,n+2} \supset \ldots$, and so the sequence $K_{n,n}, K_{n,n+1}, K_{n,n+2}, \ldots$ is a decreasing sequence of non-empty finite-dimensional affine subspaces. This sequence therefore stabilises at some non-empty affine subspace J_n of L_n. The identity (3.1) also implies that whenever $n \leq n' \leq m$ we have $\pi_{n,n'}(K_{n',m}) \subset K_{n,m}$, and so by taking m sufficiently large we see in particular that $\pi_{n,n'}(J_{n'}) \subset J_n$. We claim that in fact

$$\pi_{n,n'}(J_{n'}) = J_n. \tag{3.2}$$

Indeed, given $u \in J_n$, let m be sufficiently large that $J_n = K_{n,m}$ and $J_{n'} = K_{n',m}$. By definition of $K_{n,m}$, there is some $v \in L_m$ such that $u = \pi_{n,m}(v)$, and then yet another application of (3.1) then shows that

$$u = \pi_{n,n'}(\pi_{n',m}(v)). \tag{3.3}$$

However, $\pi_{n',m}(v) \in K_{n',m} = J_{n'}$ by definition of $K_{n',m}$, and so (3.3) implies that $u \in \pi_{n,n'}(J_{n'})$. Since, $u \in J_n$ was arbitrary, this proves (3.2), as claimed.

We now construct recursively a sequence of functions $w_n \in J_n$, $n \in \mathbb{N}$, as follows. Initially, choose an arbitrary function $w_1 \in J_1$. Then, given $w_n \in J_n$, choose w_{n+1} arbitrarily from the set $\pi_{n,n+1}^{-1}(w_n) \subset J_{n+1}$, which is non-empty by (3.2). Since w_{n+1} and w_n agree on $B(n)$, there exists $w \in V^\Gamma$ such that $w|_{B(n)} = w_n$ for every n. However, $\tau(w)|_{B(n-1)} = \tau_n(w_n) = f|_{B(n-1)}$ for every n by construction, and so $\tau(w) = f$. □

Proof of Theorem 1.9. A locally specifiable map is pre-injective on Γ if and only if it is pre-injective on every connected component of Γ, and surjective on Γ if and only if it is surjective on every connected component of Γ, and so we may assume that Γ is connected. This is essentially the same

as a reduction to the countable case made by Ceccherini-Silberstein and Coornaert in their original proof of Theorem 1.8 [7].

We first prove that surjectivity of τ implies pre-injectivity of τ'. Given $v, w \in V = \mathbb{K}^r$, write $v \cdot w = \sum_{i=1}^r v_i w_i$, and given $f_1 \in V_0^\Gamma$ and $f_2 \in V^\Gamma$ write $f_1 \cdot f_2 = \sum_{x\in\Gamma}(f_1(x) \cdot f_2(x))$. Then if τ is surjective and $\varphi \in V_0^\Gamma$, we have

$$\begin{aligned}\tau'(\varphi) = 0 &\Rightarrow \tau'(\varphi) \cdot f = 0 \text{ for every } f \in V^\Gamma \\ &\Rightarrow \varphi \cdot \tau(f) = 0 \text{ for every } f \in V^\Gamma \\ &\Rightarrow \varphi = 0\end{aligned}$$

by surjectivity of τ, and so τ' is pre-injective.

We now prove the harder direction, namely that pre-injectivity of τ' implies surjectivity of τ. Lemma 3.1 means that in order to prove that τ is surjective it suffices to show that the linear map $\tau_n : V^{B(n)} \to V^{B(n-1)}$ induced by τ is surjective. Since τ_n is a map between finite-dimensional spaces, it therefore suffices to show that its dual $\tau_n^* : V^{B(n-1)} \to V^{B(n)}$ is injective. However, the matrix of τ_n^* is precisely τ' restricted to $V^{B(n-1)}$ in domain and $V^{B(n)}$ in range, and so pre-injectectivity of τ' implies injectivity of τ_n^*, which in turn implies surjectivity of τ_n, as required. □

4. Transpose-harmonic functions and surjectivity of Laplacians

In this section we prove Propositions 1.6 and 1.7. The proofs essentially consist of a fairly direct applications of Theorem 1.9.

Definition 4.1 (Transpose-harmonic function)**.** Given a Laplacian Δ on a graph or a group Γ, we denote by Δ' the transpose of Δ, and say that a function $h : \Gamma \to \mathbb{R}$ is *transpose harmonic* if $\Delta' h = 0$.

If $\Delta = \Delta_\mu$ is the Laplacian on a group defined by a finitely supported generating probability measure μ then, writing μ' for the finitely supported generating probability measure defined by $\mu'(g) = \mu(g^{-1})$ we have

$$(\Delta_\mu)' = \Delta_{\mu'}. \tag{4.1}$$

In the case of the Laplacian on a weighted graph, on the other hand, we have the following.

Lemma 4.2. *Let Δ be the Laplacian on a locally finite weighted graph Γ, and let $f : \Gamma \to \mathbb{R}$ be a function. Then for each $x \in \Gamma$ we have*

$$\Delta' f(x) = f(x) - \sum_{y\sim x} \frac{\omega_{xy} f(y)}{\deg y}$$

In particular, f is transpose harmonic at x if and only if the function $\hat{f} : \Gamma \to \mathbb{R}$ defined by

$$\hat{f}(y) = \frac{f(y)}{\deg y}$$

is harmonic at x.

Proof. The matrix of Δ is not hard to describe. In the row corresponding to the point x, the matrix has 1 in the column corresponding to x; it has $-\omega_{xy}/\deg x$ in each column corresponding to a neighbour y of x; and every other entry is zero. The x row in the matrix of Δ' therefore has 1 in the column corresponding to x; for each neighbour y of x it has $-\omega_{xy}/\deg y$ in the column corresponding to y; and every other entry is zero. The desired result follows immediately. □

Proof of Propositions 1.6 and 1.7. In each case, Theorem 1.9 shows that it is sufficient to prove that a finitely supported transpose-harmonic function is identically zero.

In the case of the Laplacian on an infinite, connected, locally finite weighted graph (as in Proposition 1.6), Lemma 4.2 implies that the required statement is equivalent to showing that a finitely supported harmonic function is identically zero, since $\hat{f}(x) = 0$ if and only if $f(x) = 0$.

In the case of the Laplacian defined by a finitely supported generating probability measure μ (as in Proposition 1.7), (4.1) implies that the required statement is equivalent to showing that a finitely supported μ'-harmonic function is identically zero.

In each case, the required statement follows from the maximum principle (Lemma 2.2), and so the propositions are both proved. □

Remarks 4.3. The proof just presented is modelled on the amenable case of the proof of [6, Theorem 1.1], which is Proposition 1.7 in the special case that μ is uniform on a finite symmetric generating set. The proof of [6, Theorem 1.1] in the amenable case uses Theorem 1.8 in place of Theorem 1.9. The fact that Theorem 1.8 does not necessarily hold in non-amenable groups forces the authors to use a different argument in that case, in particular relying on a spectral criterion for amenability of finitely generated groups due to Kesten and Day. Our use of Theorem 1.9 allows us to avoid this complication.

Our arguments would also prove Proposition 1.6 for an *asymetrically weighted graph*, which is to say if we were to drop the assumption that $\omega_{xy} = \omega_{yx}$, provided it satisfied $\sum_{y\sim x} \omega_{xy} = \sum_{y\sim x} \omega_{yx}$ for every x.

5. A duality result for harmonic functions

The aim of this section is to prove the following result.

Proposition 5.1 (Duality result for harmonic functions)**.** *Let Γ be an infinite, connected, locally finite weighted graph, and let X be a finite subset of Γ. Then the following statements are equivalent.*

1. *Every function $f : X \to \mathbb{R}$ extends to a harmonic function on all of Γ.*
2. *There is no non-zero finitely supported function on Γ that is harmonic on $\Gamma\backslash X$.*

Remark 5.2. Proposition 5.1 fails in a finite graph, or a graph with a finite connected component, since statement (2) never holds in a finite graph, but statement (1) holds in an arbitrary graph when X is a singleton. See Remark 5.5 for details on where the proof breaks down.

Given a subset Y of Γ, we denote by $\mathbb{R}^\Gamma_Y$ the subspace of $\mathbb{R}^\Gamma$ consisting of those functions supported on Y. Proposition 5.1 then follows from combining the following two lemmas with Proposition 1.6, which implies that $\Delta(\mathbb{R}^\Gamma) = \mathbb{R}^\Gamma$.

Lemma 5.3. *Let Γ be a locally finite weighted graph, and let $X \subset \Gamma$ be a finite set. Then the following statements are equivalent.*

1. *We have $\Delta(\mathbb{R}^\Gamma_{\Gamma\backslash X}) = \mathbb{R}^\Gamma$.*
2. *There is no non-zero finitely supported function on Γ that is harmonic on $\Gamma\backslash X$.*

Lemma 5.4. *Let Γ be a locally finite weighted graph, and let $X \subset \Gamma$ be a finite set. Then the following statements are equivalent.*

1. *We have $\Delta(\mathbb{R}^\Gamma_{\Gamma\backslash X}) = \Delta(\mathbb{R}^\Gamma)$.*
2. *Every function $f : X \to \mathbb{R}$ extends to a harmonic function on all of Γ.*

Proof of Lemma 5.3. First note that by Lemma 4.2 and the fact that for every function $f : \Gamma \to \mathbb{R}$ we have $\hat{f}(x) = 0$ if and only if $f(x) = 0$, statement (2) of Lemma 5.3 is equivalent to the following statement.

(2′) There is no non-zero finitely supported function on Γ that is transpose harmonic on $\Gamma\backslash X$.

Abusing notation slightly, we identify the operator Δ with its (possibly infinite) matrix. Statement (1) of the lemma is then equivalent to saying that the matrix $\Delta_{\Gamma\backslash X}$ obtained by replacing the columns of Δ corresponding to the elements of X with columns of zeros is surjective.

Statement (2′), on the other hand, means that if $f \in \mathbb{R}_0^\Gamma$ is non-zero then $\Delta'(f)$ cannot be zero on $\Gamma\backslash X$. Put another way, this says that even if we replace the rows of Δ' corresponding to the elements of X with columns of zeros then Δ' will be pre-injective.

However, Δ' with the rows corresponding to X replaced by zeros is equal to the transpose of $\Delta_{\Gamma\backslash X}$. Replacing some entries of Δ by zeros does not change the fact that it is a locally specifiable map, and so the equivalence of (1) and (2′) therefore follows from Theorem 1.9. □

Proof of Lemma 5.4. We first prove that (1) implies (2). Let $f : X \to \mathbb{R}$ be arbitrary, and define $\overline{f}$ to be the function on Γ that agrees with f on X and takes the value 0 elsewhere. By (1) we can find a function h supported on $\Gamma\backslash X$ such that $\Delta(h) = \Delta(-\overline{f})$. The function $h + \overline{f}$ is then a harmonic extension of f, and so (2) is proved.

Conversely, note that in order to prove (1) it suffices to prove that for every $x \in X$ the function $\Delta(\delta_x)$ lies in the space $\Delta(\mathbb{R}^\Gamma_{\Gamma\backslash X})$. However, if we assume (2) then in particular we have a harmonic extension h of the function $f : X \to \mathbb{R}$ taking the value 1 at x and 0 on $X\backslash\{x\}$, and it immediately follows that $\Delta(\delta_x) = \Delta(-h|_{\Gamma\backslash X})$. □

Remark 5.5. In the case that Γ has a finite connected component, Proposition 1.6 no longer holds, and so Lemmas 5.3 and 5.4 no longer combine to prove Proposition 5.1.

6. Existence of non-constant harmonic functions on graphs

In this section we use Proposition 5.1 to prove the following result, which generalises Proposition 1.4 in the transient case.

Proposition 6.1. *Let Γ be a locally finite vertex-transitive weighted graph, and suppose that the random walk on Γ is transient. Suppose that K is finitely generated subgroup of $\operatorname{Aut}\Gamma$ such that the orbit Ke is infinite. Then there exists a harmonic function on Γ that is not constant on Ke.*

Remarks 6.2. Proposition 6.1 applies in particular to groups with symmetric, finitely supported generating probability measures, since they can be realised as vertex-transitive weighted graphs by considering their weighted Cayley graphs.

Proposition 6.1 does not necessarily hold if K has finite orbits. For example, if $G = \mathbb{Z}^3 \oplus \mathbb{Z}/2\mathbb{Z}$ and $S = \{(\pm e_1, 0), (\pm e_1, 1), (\pm e_2, 0), (\pm e_2, 1), (\pm e_3, 0), (\pm e_3, 1), (0, 1)\}$ and Γ is the Cayley graph (G, S), then every harmonic function on G is constant on the orbits of $\mathbb{Z}/2\mathbb{Z}$.

Let us note how Proposition 6.1 implies the transient case of Proposition 1.4. Proposition 1.4 is trivial when Γ is not connected; when Γ is connected and transient it follows immediately from Proposition 6.1 and the following lemma.

Lemma 6.3. *Let Γ be a connected, locally finite, vertex-transitive weighted graph. Then there is a finitely generated subgroup $G < \operatorname{Aut}\Gamma$ that is transitive.*

Proof. Let $e \in \Gamma$. By the transitivity of $\operatorname{Aut}\Gamma$, for each neighbour y of e there is an automorphism g_y of Γ such that $g_y e = y$. We claim that $G := \langle\, g_y : y \sim e \,\rangle$ is transitive; since Γ is locally finite, this is sufficient to prove the lemma.

Since Γ is connected, it suffices to show that if $z \in Ge$ and $x \sim z$ then $x \in Ge$. To see this, note that for $z \in Ge$ there exists $h \in G$ such that $e = hz$. However, this means that we have $hx \sim e$, and so $x = h^{-1} g_{hx} e \in Ge$, as desired, and the lemma is proved. □

We also recover from Proposition 6.1 the following well-known fact.

Corollary 6.4. *Let G be an infinite group with a symmetric, finitely supported generating probability measure μ. Then (G, μ) admits a non-constant harmonic function.*

Proof. If the random walk on (G, μ) is transient then the corollary follows immediately from Proposition 6.1 if we let G act on its own Cayley graph by left multiplication and take $K = G$. If the random walk is recurrent then Proposition 2.12 implies that G has either $\mathbb{Z}$ or $\mathbb{Z}^2$ as a finite-index subgroup, in which case the corollary follows from [23] or from Lemma 7.2, below. □

For the remainder of this section we are concerned with proving Proposition 6.1. Throughout, Γ is a locally finite vertex-transitive weighted graph with distinguished vertex e.

By Proposition 5.1, in order to prove Proposition 6.1 in the connected case it suffices to find two points $x, y \in Ke$ with the property that there is no non-zero finitely supported function on Γ that is harmonic except at x, y. The following result gives a necessary condition for the existence of such a function.

Lemma 6.5. *Let $x, y \in \Gamma$ and suppose that there exists a finitely-supported non-zero function $f : \Gamma \to \mathbb{R}$ that is harmonic except at x and y. Then there exists some $N > 0$ such that the conditional probability $\mathbb{P}_g[\, T_x < T_y \,|\, \min\{T_x, T_y\} < \infty \,]$ is independent of g for $d(e, g) \geq N$.*

Proof. Since f is finitely supported, there is some $N > d(e, x), d(e, y)$ such that $f(g) = 0$ whenever $d(e, g) \geq N$. We prove that the lemma holds with this N.

For $M \in \mathbb{N}$ we denote by $B(M) = B_e(M)$ the ball of radius M about the vertex e, and by τ_M the quantity $\tau_M = \min\{t : X_t \in (G \backslash B(M)) \cup \{x, y\}\}$, where $X_0, X_1, \dots$ is, as usual, the random walk on Γ. By countable additivity of $\mathbb{P}$, for $g \in B(M)$ we have

$$\mathbb{P}_g[\, X_{\tau_M} = x \,|\, X_{\tau_M} \in \{x, y\} \,] \to \mathbb{P}_g[\, T_x < T_y \,|\, \min\{T_x, T_y\} < \infty \,], \tag{6.1}$$

$$\mathbb{P}_g[\, X_{\tau_M} = y \,|\, X_{\tau_M} \in \{x, y\} \,] \to \mathbb{P}_g[\, T_y < T_x \,|\, \min\{T_x, T_y\} < \infty \,] \tag{6.2}$$

as $M \to \infty$.

Let $M \geq N$. By Lemma 2.3, there is a unique function $f_M : B(M+1) \to \mathbb{R}$ that is harmonic on $B(M) \backslash \{x, y\}$ and satisfies the following conditions:

$$f_M(x) = f(x); \tag{6.3}$$

$$f_M(y) = f(y); \tag{6.4}$$

$$f_M(z) = 0 \text{ for } z \notin B(M); \tag{6.5}$$

indeed, Lemma 2.3 implies that

$$f_M(g) = f(x) \cdot \mathbb{P}_g[\, X_{\tau_M} = x \,] + f(y) \cdot \mathbb{P}_g[\, X_{\tau_M} = y \,] \tag{6.6}$$

for $g \in B(M)$.

The restriction $f|_{B(M+1)}$ is of course harmonic on $B(M)\backslash\{x,y\}$, and trivially satisfies (6.3) and (6.4); by the definitions of N and M it also satisfies condition (6.5), and so by the uniqueness of f_M it follows that

$$f_M = f|_{B(M+1)}. \tag{6.7}$$

By the maximum principle (Lemma 2.2), and since f is not identically zero, f must be non-zero at at least one of x and y; without loss of generality we may therefore assume that $f(x) \neq 0$. If $N \leq |g| \leq M$ then (6.7) and the definition of N together imply that $f_M(g) = 0$, and so (6.6) implies that

$$\frac{\mathbb{P}_g[\,X_{\tau_M} = x\,]}{\mathbb{P}_g[\,X_{\tau_M} = y\,]} = -\frac{f(y)}{f(x)},$$

and hence that

$$\frac{\mathbb{P}_g[\,X_{\tau_M} = x \,|\, X_{\tau_M} \in \{x,y\}\,]}{\mathbb{P}_g[\,X_{\tau_M} = y \,|\, X_{\tau_M} \in \{x,y\}\,]} = -\frac{f(y)}{f(x)}.$$

Letting $M \to \infty$, we therefore see from (6.1) and (6.2) that

$$\frac{\mathbb{P}_g[\,T_x < T_y \,|\, \min\{T_x,T_y\} < \infty\,]}{\mathbb{P}_g[\,T_y < T_x \,|\, \min\{T_x,T_y\} < \infty\,]} = -\frac{f(y)}{f(x)}. \tag{6.8}$$

Since the numerator and denominator of the left-hand side of (6.8) always sum to 1, this determines $\mathbb{P}_g[\,T_x < T_y \,|\, \min\{T_x,T_y\} < \infty\,]$ uniquely and independently of g, and so the lemma is proved. □

The following lemma proves the intuitively reasonable result that if the random walk is more likely to hit x than y eventually, then it is also more likely to hit x first.

Lemma 6.6. *If $x, y \in \Gamma$ satisfy*

$$\mathbb{P}_e[\,T_x < \infty\,] > \mathbb{P}_e[\,T_y < \infty\,] \tag{6.9}$$

then they also satisfy

$$\mathbb{P}_e[\,T_x < T_y \,|\, \min\{T_x,T_y\} < \infty\,] > 1/2. \tag{6.10}$$

If the random walk on Γ is transient then (6.9) and (6.10) are equivalent.

Remark 6.7. The conditions (6.9) and (6.10) are not necessarily equivalent in a vertex-transitive graph with a recurrent random walk, as can be seen by setting $e = 0$, $x = 1$ and $y = 2$ in the Cayley graph $(\mathbb{Z}, \{\pm 1\})$.

Proof of Lemma 6.6. Write $p(x,y) = \mathbb{P}_x[\,T_y < \infty\,]$, the probability that the random walk starting at x hits y eventually. If $\mathbb{P}_e[\,T_x < \infty\,] > \mathbb{P}_e[\,T_y < \infty\,]$ then this implies in particular that $\mathbb{P}_e[\,T_z < \infty\,]$ is not constant in z, which implies that the random walk is transient. We may therefore assume that the random walk is transient and prove that (6.9) and (6.10) are equivalent.

Write

$$p(x) = \mathbb{P}_e[\,T_x < \infty \,|\, \min\{T_x,T_y\} < \infty\,],$$
$$p(y) = \mathbb{P}_e[\,T_y < \infty \,|\, \min\{T_x,T_y\} < \infty\,],$$

and note that condition (6.9) is equivalent to $p(x) > p(y)$. Write $f(x) = \mathbb{P}_e[\,T_x < T_y \,|\, \min\{T_x,T_y\} < \infty\,]$ and $f(y) = \mathbb{P}_e[\,T_y < T_x \,|\, \min\{T_x,T_y\} < \infty\,]$. Condition (6.10) is that $f(x) > 1/2$, or equivalently that $f(x) > f(y)$, since $f(x) + f(y) = 1$. However, we have $p(y) = f(y) + f(x)p(x,y)$, and by Proposition 2.7 we have $p(x,y) = p(y,x)$, and hence $p(x) = f(x) + f(y)p(x,y)$. The equivalene of (6.9) and (6.10) therefore follows, since transience of the random walk and symmetry of p together imply that $p(x,y) < 1$. □

Proposition 6.8. *If the random walk on Γ is transient then $\mathbb{P}_x[\,T_y < \infty\,] \to 0$ as $d(x,y) \to \infty$.*

Proof. It is clear that $\mathbb{P}_x[\,T_y < \infty\,] \leq \sum_{n=0}^{\infty} \mathbb{P}_x[\,X_n = y\,]$. However, since $\mathbb{P}_x[\,X_n = y\,] = 0$ for $n < d(x,y)$, we in fact have the stronger bound $\mathbb{P}_x[\,T_y < \infty\,] \leq \sum_{n=d(x,y)}^{\infty} \mathbb{P}_x[\,X_n = y\,]$. If n is even then we have $\mathbb{P}_x[\,X_n = y\,] \leq \mathbb{P}_e[\,X_n = e\,]$ by Lemma 2.5. If n is odd, on the other hand, then we have $\mathbb{P}_x[\,X_n = y\,] = \mathbb{E}_{s\in S}\mathbb{P}_{xs}[\,X_{n-1} = y\,] \leq \mathbb{P}_e[\,X_{n-1} = e\,]$, again by Lemma 2.5. Combining these last three inequalities shows that

$$\mathbb{P}_x[\,T_y < \infty\,] \leq 2 \sum_{\substack{n \geq d(x,y)-1 \\ n \text{ even}}} \mathbb{P}_e[\,X_n = e\,] \tag{6.11}$$

Recall that R_e is the number of times the random walks hits the vertex e. In particular, $R_e = \sum_{n=0}^{\infty} 1_{\{X_n = e\}}$, and so by linearity of expectation we have $\mathbb{E}_e[R_e] = \sum_{n=0}^{\infty} \mathbb{P}_e[\,X_n = e\,]$. Lemma 2.11 therefore implies that $\sum_{n=0}^{\infty} \mathbb{P}_e[\,X_n = e\,] < \infty$, which, combined with (6.11), shows that $\mathbb{P}_x[\,T_y < \infty\,] \to 0$ as $d(x,y) \to \infty$, as desired. □

Proof of Proposition 6.1. If the orbit Ke has non-trivial intersection with two connected components of Γ then the result follows by taking a function that takes the value 1 on one of these components and 0 elsewhere on Γ. We may therefore assume that Γ is connected, and so by Proposition 5.1 it suffices to find two points $x, y \in Ke$ with the property that there is no non-zero finitely supported function on Γ that is harmonic except at x, y.

We consider the following two cases.

(1) The subgroup K contains an element v such that the vertices $v^n e$ are all distinct for $n \in \mathbb{N}$.
(2) For every element u of the subgroup K there is some m such that $u^m e = e$.

In case (1), Proposition 6.8 implies that $\mathbb{P}_e[\,T_{v^n e} < \infty\,] \to 0$ and $\mathbb{P}_e[\,T_{v^{-n} e} < \infty\,] \to 0$ as $n \to \infty$. This implies that there are infinite increasing sequences $n_1^+, n_2^+, n_3^+, \ldots$ and $n_1^-, n_2^-, n_3^-, \ldots$ such that $\mathbb{P}_e[\,T_{v^{n_i^+} e} < \infty\,] > \mathbb{P}_e[\,T_{v^{n_i^+} ve} < \infty\,]$ and $\mathbb{P}_e[\,T_{v^{-n_i^-} e} < \infty\,] < \mathbb{P}_e[\,T_{v^{-n_i^-} ve} < \infty\,]$, which by Lemma 6.6 means that

$$\begin{aligned} \mathbb{P}_{v^{-n_i^+} e}[\,T_e > T_{ve} \mid \min\{T_e, T_{ve}\} < \infty\,] &> 1/2, \\ \mathbb{P}_{v^{n_i^-} e}[\,T_e > T_{ve} \mid \min\{T_e, T_{ve}\} < \infty\,] &< 1/2. \end{aligned}$$

Since $v^{-n_i^+} e \to \infty$ and $v^{n_i^-} e \to \infty$, Lemma 6.5 therefore implies that there exists no finitely supported non-zero function on Γ that is harmonic except at e, ve, and so the proposition is proved in case (1).

In case (2), let R be a finite symmetric generating set for K. We claim that there are elements $x_1, x_2, \ldots \in K$ with $d(e, x_n e) \to \infty$ such that, for each n, there is some $r_n \in R$ such that $\mathbb{P}_e[\,T_{x_n e} < \infty\,] < \mathbb{P}_e[\,T_{x_n r_n e} < \infty\,]$. Indeed, for each $n = 1, 2, \ldots$, let x_n be a point of minimal distance from the identity in the Cayley graph (K, R) such that $\mathbb{P}_e[\,T_{x_n e} < \infty\,] < 1/n$. Such a point always exists by Proposition 6.8 and the assumption that the orbit Ke is infinite, and by the regularity and local finiteness of Γ we have

$$d(e, x_n e) \to \infty \tag{6.12}$$

as $n \to \infty$. By definition of x_n, and using (6.12), for sufficiently large n there is some $r_n \in R$ such that $\mathbb{P}_e[\,T_{x_n r_n e} < \infty\,] \geq 1/n > \mathbb{P}_e[\,T_{x_n e} < \infty\,]$, as caimed.

By the finiteness of R, upon passing to a subsequence if necessary we may in fact assume that there is some $u \in R$ such that for each n we have

$$\mathbb{P}_e[\,T_{x_n e} < \infty\,] < \mathbb{P}_e[\,T_{x_n u e} < \infty\,]. \tag{6.13}$$

We claim that there is no non-zero finitely supported function on Γ that is harmonic except at e, ue.

As in case (1), condition (6.13) and Lemma 6.6 imply that

$$\mathbb{P}_{x_n^{-1}e}\left[\, T_e > T_{ue} \,\middle|\, \min\{T_e, T_{ue}\} < \infty \,\right] \;<\; 1/2;$$

indeed, applying the automorphism u^m, we see that

$$\mathbb{P}_{u^m x_n^{-1}e}\left[\, T_{u^m e} > T_{u^{m+1}e} \,\middle|\, \min\{T_{u^m e}, T_{u^{m+1}e}\} < \infty \,\right] \;<\; 1/2 \tag{6.14}$$

for every $m \in \mathbb{N}$. Moreover, (6.12) implies that for each $m \in \mathbb{N}$ we have $d(u^{-m}e, x_n^{-1}e) \to \infty$ as $n \to \infty$, and so $d(e, u^m x_n^{-1}e) \to \infty$ as $n \to \infty$. If the claim is false, and there does exist some non-zero finitely supported function on Γ that is harmonic except at e, ue, then translating this function by u^m we see that there is also a function on Γ that is harmonic except at $u^m e, u^{m+1}e$. Combining (6.12) and (6.14) with Lemma 6.5 therefore implies that for each m there is some $N_m > 0$ such that

$$\mathbb{P}_x\left[\, T_{u^m e} > T_{u^{m+1}e} \,\middle|\, \min\{T_{u^m e}, T_{u^{m+1}e}\} < \infty \,\right] \;<\; 1/2$$

for every $x \in \Gamma$ such that $d(e, x) \geq N_m$; since the orbit of e under u is finite we may assume that the N_m are all equal to some $N > 0$. Fixing some x with $d(e, x) \geq N$ and applying Lemma 6.6 once more, this means that

$$\mathbb{P}_x[\, T_{u^m e} < \infty \,] < \mathbb{P}_x[\, T_{u^{m+1}e} < \infty \,]$$

for every $m \in \mathbb{N}$, which implies by induction that

$$\mathbb{P}_x[\, T_e < \infty \,] < \mathbb{P}_x[\, T_{u^m e} < \infty \,]$$

for every $m \in \mathbb{N}$. This is impossible, however, since there is some $m \in \mathbb{N}$ such that $u^m e = e$, and so it must have been the case that there was no non-zero finitely supported function on Γ harmonic except at e, ue. This proves the claim, and hence the proposition in case (2). □

7. Harmonic functions on virtually abelian groups

In this section we investigate spaces of harmonic functions on virtually abelian groups. The first purpose is to prove the easier direction of Theorem 1.1, as follows.

Proposition 7.1 (Direct statement of Theorem 1.1)**.** *Let G be a group with a finite-index subgroup isomorphic to $(\mathbb{Z}, +)$, and let μ be a symmetric, finitely supported generating probability measure on G. Then* $\dim H(G, \mu) < \infty$.

The second is to note a characterisation of the space $H^1(G, \mu)$ of harmonic functions of linear growth on a virtually abelian group G (see Lemma 7.2, below).

Let G be a group with a finite-index normal subgroup isomorphic to $\mathbb{Z}^d$, and let μ be a symmetric, finitely supported generating probability measure on G. Abbreviate $S := \operatorname{supp}\mu$. Fix a right-transversal T of $\mathbb{Z}^d$ containing the identity, which is to say a finite set T such that each $g \in G$ can be expressed uniquely as $g = \zeta(g)\tau(g)$ with $\zeta(g) \in \mathbb{Z}^d$ and $\tau(g) \in T$. We write $\zeta_i(g)$ for the ith coordinate of $\zeta(g)$ with respect to the standard basis for $\mathbb{Z}^d$.

Lemma 7.2 ([2, 23])**.** *For each $i = 1, \ldots, d$ there is a function $\varphi_i : G \to \mathbb{R}$ that factors through $G/\mathbb{Z}^d$ such that the function $f_i : G \to \mathbb{R}$ given by $f_i(g) = \zeta_i(g) + \varphi_i(\tau(g))$ is harmonic on (G, μ). Moreover, $H^1(G, \mu)$ is spanned by the set $\{1, f_1, \ldots, f_d\}$.*

Proof. The existence of the harmonic functions f_i follows directly from [23, Theorem 3.6]. The fact that $\{1, f_1, \ldots, f_d\}$ spans $H^1(G, \mu)$ is then precisely the linear-growth case of [2, Theorem 1.12]; see also [20] for a more elementary proof. □

Lemma 7.3. *Let $d_{\mathbb{Z}^d}$ be the Cayley-graph distance on $\mathbb{Z}^d$ with respect to the standard generating set. Then there exists $M \in \mathbb{N}$ such that for every $g \in G$ and every $s \in S$ we have $d_{\mathbb{Z}^d}(\zeta(gs), \zeta(g)) \leq M$.*

Proof. Given $g \in G$ and $s \in S$, write $t = \tau(g)$, so that $gs = \zeta(g)ts = \zeta(g)\zeta(ts)\tau(ts)$. This implies, in particular, that $\zeta(gs) = \zeta(g)\zeta(ts)$, and so we may take M to be the maximum over the (finite) set $\{|\zeta(ts)|_{\mathbb{Z}^d} : s \in S, t \in T\}$. □

Proof of Proposition 7.1. Lemma 7.3 implies that for each $n \in \mathbb{Z}$ we have $[-n, n]TS \subset [-n-M, n+M]T$. It follows that $([-n, n]T)^+ \subset [-n-M, n+M]T$, and so $\partial^+([-n, n]T)$ has cardinality at most $2M|T|$. Lemma 2.3 therefore implies that the space of functions on $\partial^+([-n, n]T)$ that are harmonic on $[-n, n]T$ is of dimension at most $2M|T|$. However, $G = \bigcup_{n=1}^{\infty}[-n, n]T$, and so the space of harmonic functions on G is also of dimension at most $2M|T|$. □

Remark 7.4. Taking $G = \mathbb{Z}$ and setting μ to be the uniform probability measure on $[-M, M]$ shows that the bound $2M|T|$ on the dimension of the space of harmonic functions in the proof of Proposition 7.1 can be tight. In particular, the precise dimension depends on the measure μ as well as on the group G.

8. Positive harmonic functions on linear groups

If G is a group and μ is a finitely supported generating probability measure then a *positive harmonic function* on (G, μ) is a harmonic function $h : G \to \mathbb{R}$ that takes only positive values. G. Margulis [18] showed that a nilpotent group admits no non-constant positive harmonic functions. More generally, we have the following result of W. Hebish and L. Saloff Coste.

Proposition 8.1 (Hebish–Saloff Coste [15])**.** *Let G be a virtually nilpotent group with a symmetric, finitely supported generating probability measure μ. Then (G, μ) admits no non-constant positive harmonic functions.*

P. Bougerol and L. Elie show that for linear groups the converse is also true.

Proposition 8.2 (Bougerol–Elie [4])**.** *Let G be a subgroup of $GL_d(\mathbb{R})$ that is not virtually nilpotent, and let μ be a symmetric, finitely supported generating probability measure on G. Then (G, μ) admits a non-constant positive harmonic function.*

The purpose of this section is to show that, in that case, there are in fact many positive harmonic functions.

Proposition 8.3. *Let G be a group with a symmetric, finitely supported generating probability measure μ, and suppose that (G, μ) admits at least one non-constant positive harmonic function. Then the set of positive harmonic functions on (G, μ) spans an infinite-dimensional space.*

The following is then immediate.

Corollary 8.4. *Let G be a subgroup of $GL_d(\mathbb{R})$ that is not virtually nilpotent, and let μ be a symmetric, finitely supported generating probability measure on G. Then the positive harmonic functions on (G, μ) span an infinite-dimensional space.*

Question 8.5. Does an arbitrary non-virtually nilpotent group with a symmetric, finitely supported generating probability measure admit a non-constant positive harmonic function?

In proving Proposition 8.3 we make use of the *minimal Martin boundary* of (G, μ).

Definition 8.6 (Minimal harmonic function)**.** Given a group G with a finitely supported generating probability measure μ, a *minimal harmonic function* on (G, μ) is a positive harmonic function $f : G \to \mathbb{R}$ with the property that every other positive harmonic function $f' : G \to \mathbb{R}$ satisfying $f' \leq f$ is a constant multiple of f. A *normed minimal harmonic function* $f : G \to \mathbb{R}$ is a minimal harmonic function satisfying $f(e) = 1$.

Definition 8.7 (Minimal Martin boundary)**.** The *minimal Martin boundary* $\Delta(G, \mu)$ of the pair (G, μ) is the compact closure, in the topology of pointwise convergence, of the set of normed minimal harmonic functions on (G, μ).

Each positive harmonic function $f : G \to \mathbb{R}$ has a unique *representing measure* ν_f on the minimal Martin boundary, which is to say a measure ν_f such that

$$f(x) = \int_\Delta h(x) d\nu_f(h) \tag{8.1}$$

for every $x \in G$ (see [16, §0.3] or [26, §7, p. 32]).

Lemma 8.8. *The set of normed minimal harmonic functions on a group G with respect to a finitely supported generating probability measure μ is linearly independent.*

Proof. Suppose that $h_1, \ldots, h_r$ are distinct minimal harmonic functions and let $\alpha_1, \ldots, \alpha_r$ be such that $\sum_{i=1}^m \alpha_i h_i = 0$. Without loss of generality we may assume that $\alpha_i \leq 0$ for $i \leq k$, and that $\alpha_i \geq 0$ for $i > k$, and so in fact we have $\sum_{i=1}^k (-\alpha_i) h_i = \sum_{i=k+1}^m \alpha_i h_i$. However, both the left-hand side and the right-hand side of this expression are non-negative harmonic functions, and so it follows from the uniquness of the representation (8.1) that the α_i are all zero. □

Proof of Proposition 8.3. We prove the contrapositive. Suppose that the set of positive harmonic functions on G does not span an infinite-dimensional space. By Lemma 8.8 this implies in particular that the set of normed minimal harmonic functions is finite, so we may enumerate them as $h_1, \ldots, h_m$.

The group G acts on the space of all harmonic functions via $g \cdot f(x) = f(g^{-1}x)$. The image of a minimal harmonic function under this action is another minimal harmonic function, and so in particular for each $i = 1, \ldots, m$ and each $g \in G$ we have some $\alpha_{g,i} \in \mathbb{R}$ and some $g \cdot i \in [m]$ such that $g \cdot h_i = \alpha_{g,i} h_{g \cdot i}$. As the notation $g \cdot i$ implicitly suggests, this defines an action of G on the set $[m]$.

By the orbit-stabiliser theorem, for each i the stabiliser H_i of i is of finite index in G; by Lemma 2.13, we may set H to be a normal subgroup of G that has finite index in $\bigcap_{i=1}^m H_i$, and hence in G. For every $g \in H$ we have $g \cdot h_i = \alpha_{g,i} h_i$, which is to say that $h_i(g^{-1}x) = \alpha_{g,i} h_i(x)$ for every $x \in G$ and every i. Taking $x = e$, and noting that $h_i(e) = 1$, we see that $\alpha_{g,i} = h_i(g^{-1})$, and so this implies that $h_i(g^{-1}x) = h_i(g^{-1}) h_i(x)$ for every $g \in H$ and every $x \in G$.

This implies that the restriction of h_i to H is a homomorphism into $\mathbb{R}^\times$, and moreover that $h_i(cx) = h_i(x)$ for every $c \in [H, H]$ and every $x \in G$. We conclude that each h_i factors through $G/[H, H]$ (noting that $[H, H]$ is characteristic in H, and hence normal in G).

Let $p : G \to \mathbb{R}$ be a positive harmonic function. Since p can be expressed in the form (8.1), p must also factor through $G/[H, H]$, and so writing $\phi : G \to G/[H, H]$ we have $p = \hat{p} \circ \phi$, with $\hat{p} : G/[H, H] \to \mathbb{R}$ harmonic by Lemma 2.1. However, the abelian group $H/[H, H]$ is of finite index in $G/[H, H]$, and so Proposition 8.1 therefore implies that $\hat{p}$, and hence p, is constant. □

9. Random walks on virtually cyclic groups

In this section we consider an infinite group G with a finite-index normal cyclic subgroup $\mathbb{Z}$ and a symmetric, finitely supported generating probability measure μ. In a similar fashion to Section 7, we consider a finite set T such that each $g \in G$ can be expressed uniquely as $g = \zeta(g)\tau(g)$ with $\zeta(g) \in \mathbb{Z}$ and $\tau(g) \in T$.

In general we continue to denote the identity of G by e, the inverse of an element g by g^{-1}, and the composition of two group elements g, h by gh. However, when composing elements of $\mathbb{Z}$ with one another we often switch to additive notation to emphasise the integer structure. Thus, for example, we sometimes denote the identity by 0, the inverse of m by $-m$ and the composition of m and n by $m + n$, provided $m, n \in \mathbb{Z}$. This should not cause confusion since, whilst the notation for a given group element is not unique, neither is it ambiguous (in particular, we never multiply together two elements of $\mathbb{Z}$). For the avoidance of doubt, the notation 1 always represents a generating element of the subgroup $\mathbb{Z}$, and never the identity element of G.

For each $n \in \mathbb{N}$ write $T_n^+ = \min\{t \geq 0 : \zeta(X_t) \geq n\}$ and $T_n^- = \min\{t \geq 0 : \zeta(X_t) \leq n\}$, noting that these quantities are almost surely finite. The purpose of this section is then to prove the following result.

Lemma 9.1. *Let $m \in \mathbb{Z}$, and suppose that $g \in G$ with $m < \zeta(g) < m + R$. Let M be as in Lemma 7.3. Then*

$$\mathbb{P}_g\left[\, T_{m+R}^+ < T_m^- \,\right] = \frac{\zeta(g) - m}{R + M} + O\left(\frac{1}{R}\right)$$

Proof. By Lemma 7.2 there exists a function $\varphi : T \to \mathbb{R}$ such that the function $f : G \to \mathbb{R}$ given by $f(nt) = n + \varphi(t)$ is harmonic on G. Let $t_{\min} \in T$ be the point at which φ takes its minimum value, and $t_{\max} \in T$ the point at which φ takes its maximum value, and define two further harmonic functions $f^+, f^- : G \to \mathbb{R}$ by

$$f^+ = \frac{1}{R+M}\Big(f - f((m-M)t_{\min})\Big); \qquad f^- = \frac{1}{R+M}\Big(f - f((m+R+M)t_{\max})\Big) + 1.$$

Note the following properties of f^+, f^-.

(i) We have $f^-(nt) \leq 0 \leq f^+(nt)$ whenever $n \in [m - M, m]$.
(ii) We have $f^-(nt) \leq 1 \leq f^+(nt)$ whenever $n \in [m + R, m + R + M]$.

Moreover, $f^+ - f^-$ is constant and given by

$$f^+ - f^- = \frac{\varphi(t_{\max}) - \varphi(t_{\min}) + M}{R + M}, \tag{9.1}$$

and we have

$$f^+((n+1)t) - f^+(nt) = f^-((n+1)t) - f^-(nt) = \frac{1}{R+M} \tag{9.2}$$

for every $n \in \mathbb{Z}$ and every $t \in T$.

Now define $h : [m - M, m + R + M]T \to \mathbb{R}$ by setting

$$h(nt) = \begin{cases} 0 & \text{when } n \in [m - M, m] \\ 1 & \text{when } n \in [m + R, m + R + M], \end{cases}$$

and requiring that h be harmonic elsewhere. Lemma 2.3 and the definition of M imply that h is well defined by these stipulations, and moreover that

$$h(g) = \mathbb{P}_g\left[\, T_{m+R}^+ < T_m^- \,\right]. \tag{9.3}$$

Now Corollary 2.4, properties (i) and (ii) of f^+, f^- and the definition of h imply that $f^- \le h \le f^+$, and hence (i), (ii), (9.1) and (9.2) imply that

$$h(g) = \frac{\zeta(g) - m}{R + M} + O\left(\frac{1}{R}\right)$$

The desired result then follows from (9.3). □

10. Harmonic functions on groups with virtually cyclic quotients

In this section we consider groups with virtually cyclic quotients. A well-known example of a group with a genuinely cyclic quotient is the *lamplighter group*. If L is the $\mathbb{Z}/2\mathbb{Z}$-vector space of finitely supported functions $\mathbb{Z} \to \mathbb{Z}/2\mathbb{Z}$, viewed as an additive group, then the *lamplighter group* G is the semidirect product $G = \mathbb{Z} \ltimes L$ defined by the action of $m \in \mathbb{Z}$ on L given by $m \cdot f(x) = f(x-m)$. Explicitly, the group operation is defined by $(m, f) \cdot (m', f') = (m + m', f + m \cdot f')$.

I. Benjamini, G. Kozma and Yadin [3] give an explicit construction of a positive harmonic function on the lamplighter group.

Proposition 10.1 (Benjamini–Kozma–Yadin, unpublished)**.** *Let G be the lamplighter group, and let μ be a symmetric, finitely supported generating probability measure on G. Denote the random walk on the lamplighter group by (M_0, F_0), (M_1, F_1), $(M_2, F_2), \dots$. Let $\tau_r = \min\{t \ge 0 : |M_t| \ge r\}$, and define $h_r : G \to \mathbb{R}$ by $h_r(g) = \mathbb{P}_g[\, F_{\tau_r}(n) = 0$ for all $n < 0\,]$. Then $r h_r$ converges pointwise to a positive harmonic function on G.*

In order to prove Theorem 1.1, we need a slightly more general result. The purpose of this section is to show that the construction of Benjamini, Kozma and Yadin can be adapted fairly easily to obtain harmonic functions on a more general family of finitely generated groups with virtually cyclic quotients.

Proposition 10.2. *Let G be a group with a symmetric, finitely supported generating probability measure μ, and suppose that there is a homomorphism ψ from G onto an infinite virtually cyclic group such that $K = \ker \psi$ is not finitely generated. Then (G, μ) admits a positive harmonic function of at most linear growth that does not factor through G/K.*

Remark 10.3. The function we construct in proving Proposition 10.2 is positive, and so Proposition 8.3 implies that G has an infinite-dimensional space spanned by positive harmonic functions, although we do not need this to prove Theorem 1.1. It also implies that $H^1(G, \mu)$ is infinite dimensional, since if $\dim H^1(G, \mu) < \infty$ then every linearly growing harmonic function restricts to a homomorphism on some finite-index subgroup of G [20].

Remark 10.4. Meyerovitch and Yadin [19] generalise Proposition 10.1 in another direction in proving their result that finite dimensionality of $H^1(G, \mu)$ for G soluble implies that G is virtually nilpotent.

We start our proof of Proposition 10.2 by expressing G in a particularly convenient form.

Lemma 10.5. *The group G posseses an infinite cyclic subgroup $\mathbb{Z}$ such that $K\mathbb{Z}$ is normal in G, and a finite set T containing the identity such that each $g \in G$ can be expressed uniquely as*

$$g = \kappa(g)\zeta(g)\tau(g) \tag{10.1}$$

with $\kappa(g) \in K$, $\zeta(g) \in \mathbb{Z}$ and $\tau(g) \in T$. Moreover, ψ is an isomorphism on $\mathbb{Z}$ and injective on T, and each $\overline{g} \in \psi(G)$ can be expressed uniquely as

$$\overline{g} = \overline{\zeta}(\overline{g})\overline{\tau}(\overline{g}) \tag{10.2}$$

with $\overline{\zeta}(\overline{g}) \in \psi(\mathbb{Z})$ and $\overline{\tau}(\overline{g}) \in \psi(T)$.

Proof. The image $\psi(G)$ possesses an infinite cyclic subgroup $\langle z \rangle$ of finite index, and by Lemma 2.13 we may assume that $\langle z \rangle$ is normal in $\psi(G)$. Let $\overline{z} \in \psi^{-1}(z)$. The element $\overline{z}$ is of infinite order, and we denote by $\mathbb{Z}$ the infinite cyclic subgroup that it generates. Note that ψ is injective on $\mathbb{Z}$, and hence an isomorphism on $\mathbb{Z}$, as required.

Since $\psi(\mathbb{Z}) = \langle z \rangle$ is of finite index in $\psi(G)$, we may choose a finite set $\overline{T}$ containing e such that each $\overline{g} \in \psi(G)$ can be expressed uniquely in the form (10.2), with $\overline{\tau}(\overline{g}) \in \overline{T}$. For each $\overline{t} \in \overline{T}$ pick an arbitrary $t \in \psi^{-1}(\overline{T})$, and define $T = \{t : \overline{t} \in \overline{T}\}$. It immediately follows that the element $\overline{\tau}(\overline{g})$ in (10.2) belongs to $\psi(T)$, and that ψ is injective on T, as required. The injectivity of ψ on $\mathbb{Z}$ additionally implies that each $g \in G$ can be expressed uniquely in the form (10.1).

The fact that $K\mathbb{Z}$ is normal in G follows immediately from the fact that $\psi(\mathbb{Z}) = \langle z \rangle$ is normal in $\psi(G)$. □

From now on in this section ψ and K are as in Proposition 10.2, and $\mathbb{Z}$ and T are fixed as in Lemma 10.5. Note that we have $\overline{\tau}(\psi(g)) = \psi(\tau(g))$, and that if we abuse notation slightly and identify $\mathbb{Z}$ with its isomorphic image $\psi(\mathbb{Z})$ we have $\overline{\zeta}(\psi(g)) = \zeta(g)$.

As in Section 9, when composing elements of $\mathbb{Z}$ or $\psi(\mathbb{Z})$ with one another we often switch to additive notation to emphasise the integer structure.

Since K is normal, the group $\mathbb{Z}$ acts on K by conjugation. We may therefore define an automorphism $\varphi : K \to K$ by $\varphi(k) = 1k1^{-1}$. More generally, this means that $\varphi^n(k) = nkn^{-1}$. As in Section 7, we denote $S := \operatorname{supp} \mu$.

If $g = knt$ is a group element with $k \in K$, $n \in \mathbb{Z}$ and $t \in T$, then the elements adjacent to g in the Cayley graph (G, S) are the elements $gs = knts$ with $s \in S$.

Lemma 10.6. *Let $k \in K$, $n \in \mathbb{Z}$, $t \in T$ and $s \in G$. Then*

$$\kappa(knts) = k\varphi^n(\kappa(ts)); \qquad \zeta(knts) = n + \zeta(ts); \qquad \tau(knts) = \tau(ts).$$

Proof. Expressing ts in the form (10.1), we have $knts = kn\kappa(ts)\zeta(ts)\tau(ts)$, and hence $knts = k\varphi^n(\kappa(ts))n\zeta(ts)\tau(ts)$, as claimed. □

For each set $A \subset \mathbb{Z}$ define a subgroup U_A of K by

$$U_A = \langle\, \varphi^n(\kappa(ts)) : s \in S, t \in T, n \in A \,\rangle,$$

and for each $n \in \mathbb{Z}$ abbreviate by U_n the subgroup $U_n = U_{[n,\infty)}$. Lemma 10.6 implies that

$$K = U_{\mathbb{Z}}. \tag{10.3}$$

Lemma 10.7. *If K is not finitely generated then, possibly after relabelling each $n \in \mathbb{Z}$ as $-n$, we have*

$$\cdots \supsetneq U_{-2} \supsetneq U_{-1} \supsetneq U_0 \supsetneq U_1 \supsetneq U_2 \supsetneq \cdots. \tag{10.4}$$

Proof. The containments of (10.4) are immediate by definition, so we just need to prove that they are strict. We start by showing that either $U_{\{0\}} \not\subset U_{\mathbb{N}}$ or $U_{\{0\}} \not\subset U_{-\mathbb{N}}$. Indeed, suppose that

$U_{\{0\}} \subset U_{\mathbb{N}}$ and $U_{\{0\}} \subset U_{-\mathbb{N}}$, which, since $U_{\{0\}}$ is finitely generated, implies in particular that there is some $M \in \mathbb{N}$ such that

$$U_{\{0\}} \subset U_{[M]}; \tag{10.5}$$

$$U_{\{0\}} \subset U_{-[M]}. \tag{10.6}$$

Since φ is an automorphism, (10.5) also implies that $U_{\{-1\}} \subset U_{\{0\}\cup[M-1]}$, and hence by (10.5) that $U_{\{-1\}} \subset U_{[M]}$. Repeating this argument, we conclude that $U_{\{-n\}} \subset U_{[M]}$ for every $n \in \mathbb{N}$. Similarly, (10.6) implies that $U_{\{n\}} \subset U_{-[M]}$ for every $n \in \mathbb{N}$, and so in fact we have $U_{\mathbb{Z}} = U_{[-M,M]}$. By (10.3), this contradicts the assumption that K is not finitely generated, and so either $U_{\{0\}} \not\subset U_{\mathbb{N}}$ or $U_{\{0\}} \not\subset U_{-\mathbb{N}}$, as claimed. Upon relabelling each $n \in \mathbb{Z}$ by $-n$ if necessary, we may assume the former, which implies in particular that $U_0 \not\subset U_1$. Repeatedly using the fact that φ is an automorphism then yields the lemma. □

We assume from now on that $\mathbb{Z}$ is labelled in such a way that (10.4) holds.

As usual, we denote by $X_0, X_1, X_2, \dots$ the random walk on G defined by μ. In this section, we additionally denote by $\overline{X}_0, \overline{X}_1, \overline{X}_2, \dots$ the random walk on $\psi(G)$ defined by $\psi(\mu)$. Note that the projected walk $(\psi(X_t))$ is isomorphic to the random walk $(\overline{X}_t)$.

For each $n \in \mathbb{N}$, write

$$T_n^+ = \min\{t \geq 0 : \zeta(X_t) \geq n\}, \qquad T_n^- = \min\{t \geq 0 : \zeta(X_t) \leq n\},$$
$$\overline{T}_n^+ = \min\{t \geq 0 : \overline{\zeta}(\overline{X}_t) \geq n\}, \qquad \overline{T}_n^- = \min\{t \geq 0 : \overline{\zeta}(\overline{X}_t) \leq n\},$$

noting that each of these quantities is almost surely finite. Note, incidentally, that if we identify $\overline{X}_t = \psi(X_t)$ then $T_n^\pm = \overline{T}_n^\pm$.

Define $B_R = \min\{t \geq 0 : \zeta(X_u) \geq 0 \text{ for all } u \in [t, T_R^+]\}$. More generally, for each $n < R$ set $B_R^n = \min\{t \geq 0 : \zeta(X_u) \geq n \text{ for all } u \in [t, T_R^+]\}$.

Lemma 10.8. *There exist some $l > \max \zeta(TS)$ and some $\alpha \in (0,1)$ such that if $R > l$, and if g is such that $-l \leq \zeta(g) \leq 0$, then either*

$$\mathbb{P}_g[\, \zeta(X_t) \geq -l \text{ for all } t \leq T_R^+ \,] = 0$$

or

$$\mathbb{P}_g[\, \kappa(X_{B_R}) \in U_0 \,|\, \zeta(X_t) \geq -l \text{ for all } t \leq T_R^+ \,] \leq \alpha$$

Proof. Fix an element $u \in K \backslash U_0$, and for each $t \in T$ and each j satisfying $0 \leq j < \max \zeta(TS)$ fix a path

$$x_0^{j,t} = e, x_1^{j,t}, x_2^{j,t}, \dots, x_{r_{j,t}}^{j,t} = t^{-1}\varphi^{-j}(u)t$$

from e to $t^{-1}\varphi^{-j}(u)t$ in the Cayley graph (G, S), chosen so that

$$\zeta(t x_i^{j,t}) < -\max \zeta(TS) \tag{10.7}$$

for at least one i.

Let $l = (1 + \max_{j,t} r_{j,t}) \max \zeta(TS)$. Write $\gamma = \min_{s \in S} \mu(s)$, and set $\beta = \gamma^{\max_{j,t} r_{j,t}}$. Note that for each j, t there is a probability of at least β that the random walk starting at e has $x_0^{j,t}, x_1^{j,t}, \dots, x_{r_{j,t}}^{j,t}$ as an initial segment.

Write A for the set of (finite) paths p from g whose images $\zeta(p)$ in $\mathbb{Z}$ finish at R or above, but stay in the range $[-l, R-1]$ until then. If $A = \varnothing$ then $\mathbb{P}_g[\, \zeta(X_t) \geq -l \text{ for all } t \leq T_R^+ \,] = 0$ and the lemma holds, and so we may assume that $A \neq \varnothing$. For each $p \in A$, write $\overline{k}_p \overline{m}_p \overline{t}_p$ for the final position of p,

with $\overline{k}_p \in K$, $\overline{m}_p \in \mathbb{Z}$ and $\overline{t}_p \in T$; thus $\overline{m}_p \geq R$, but all earlier positions of $\zeta(p)$ are below R. Also, let σ_p be the largest final segment of p whose image in $\mathbb{Z}$ lies entirely in the non-negative integers, and let $k_p m_p t_p$ be the first position of this final segment, with $k_p \in K$, $m_p \in \mathbb{Z}$ and $t_p \in T$. Note that

$$0 \leq m_p < \max \zeta(TS). \tag{10.8}$$

Lemma 10.6 implies that $\{p \in A : k_p \in U_0\} = \{p \in A : \overline{k}_p \in U_0\}$, and so we may define $A_\in = \{p \in A : k_p \in U_0\} = \{p \in A : \overline{k}_p \in U_0\}$ and $A_\notin = A \backslash A_\in$. We claim that

$$\mathbb{P}_g(A_\notin) \gg \mathbb{P}_g(A_\in). \tag{10.9}$$

This is sufficient to prove the lemma, since the conditional probability we are aiming to bound is equal to

$$\frac{\mathbb{P}_g(A_\in)}{\mathbb{P}_g(A_\notin) + \mathbb{P}_g(A_\in)}.$$

We define a map c from $A_\in$ to the set of finite paths starting at g as follows. Given $p \in A_\in$, let $c(p)$ be the path that agrees with p up until $k_p m_p t_p$, then has positions $k_p m_p t_p x_1^{m_p,t_p}, \dots, k_p m_p t_p x_{r_{m_p,t_p}}^{m_p,t_p}$, and then continues with the same increments as the original path p had after position $k_p m_p t_p$. This is well defined by (10.8).

We claim that $c(p) \in A_\notin$ for every $p \in A_\in$. To see that $c(p) \in A$, note that

$$k_p m_p t_p x_{r_{m_p,t_p}}^{m_p,t_p} = k_p u m_p t_p. \tag{10.10}$$

This implies in particular that

$$\psi(k_p m_p t_p x_{r_{m_p,t_p}}^{m_p,t_p}) = m_p \psi(t_p) = \psi(k_p m_p t_p). \tag{10.11}$$

By definition of l, at no point between $k_p m_p t_p$ and $k_p m_p t_p x_{r_{m_p,t_p}}^{m_p,t_p}$ does $\zeta(p)$ drop below $-l$, and so it follows that $c(p) \in A$. To see, more specifically, that $c(p) \in A_\notin$, note that the definition of $k_p m_p t_p$ combines with (10.11) to imply that $\zeta(c(p))$ doesn't drop below zero after $k_p m_p t_p x_{r_{m_p,t_p}}^{m_p,t_p}$. Lemma 10.6 and (10.10) therefore imply that $\overline{k}_{c(p)}$ is in the same left coset of U_0 as $k_p u$. In particular, since $k_p \in U_0$ and $u \notin U_0$ we have $\overline{k}_{c(p)} \notin U_0$, and so $c(p) \in A_\notin$, as claimed.

The fact that $c(A_\in) \subset A_\notin$ of course implies that

$$\mathbb{P}_g(A_\notin) \geq \mathbb{P}_g(c(A_\in)). \tag{10.12}$$

We claim, moreover, that c is $O(1)$-to-one. Write $a(p)$ for the segment that was added to c to obtain $c(p)$, and note that one can, in principle at least, recover p from $c(p)$ simply by deleting the segment $a(p)$. Note that (10.7) and (10.8) combine with Lemma 10.6 and the fact (noted in the preceding paragraph) that $\zeta(c(p))$ doesn't drop below zero after $k_p m_p t_p x_{r_{m_p,t_p}}^{m_p,t_p}$ to imply that $\zeta(p)$ drops below zero for the last time at some point during $a(p)$. This means that knowledge of $c(p)$ only is sufficient to identify, to within $\max_{j,t} r_{j,t}$ positions, where in $c(p)$ the segment $a(p)$ begins. Furthermore, the increments of $a(p)$ coincide with those of one of the finitely many paths $(x_i^{j,t})$. There are therefore at most $O(1)$ possibilities for $a(p)$, given $c(p)$, and so c is $O(1)$-to-one, as claimed.

This implies, in particular, that

$$\mathbb{P}_g(c(A_\in)) \gg \sum_{p \in A_\in} \mathbb{P}_g(c(p)). \tag{10.13}$$

However, it follows from the definition of β that for every $p \in A_\in$ we have $\mathbb{P}_g(c(p)) \geq \beta \mathbb{P}_g(p)$. In combination with (10.12) and (10.13), this implies that

$$\mathbb{P}_g(A_{\notin}) \geq \mathbb{P}_g(c(A_\in)) \gg \sum_{p \in A_\in} \mathbb{P}_g(c(p)) \geq \beta \sum_{p \in A_\in} \mathbb{P}_g(p) = \beta \mathbb{P}_g(A_\in),$$

and so (10.9) holds as claimed and the lemma is proved. □

Lemma 10.9. *Let l and α be as given by Lemma 10.8. Let $n \leq 0$. Then if $R > l$, and if g is such that $n - l \leq \zeta(g) \leq n$, then either*

$$\mathbb{P}_g[\, \zeta(X_t) \geq n - l \text{ for all } t \leq T_R^+ \,] = 0$$

or

$$\mathbb{P}_g[\, \kappa(X_{B_R^n}) \in U_n \,|\, \zeta(X_t) \geq n - l \text{ for all } t \leq T_R^+ \,] \leq \alpha.$$

Proof. This follows immediately from applying Lemma 10.8 with the weighted Cayley graph (G, μ) left-translated by n. □

Lemma 10.10. *Let l and α be as given by Lemma 10.8. Let $m \in \mathbb{N}$, and suppose that $k \leq -ml$ and $R > l$. Then whenever $g \in G$ is such that $k \leq \zeta(g) < k + l$ we have either*

$$\mathbb{P}_g[\, \zeta(X_t) \geq k \text{ for all } t \leq T_R^+ \,] = 0$$

or

$$\mathbb{P}_g[\, \kappa(X_{T_R^+}) \in U_0 \,|\, \zeta(X_t) \geq k \text{ for all } t \leq T_R^+ \,] \leq \alpha^m.$$

Proof. Everthing in this lemma is conditional on the event $\{\, \zeta(X_t) \geq k \text{ for all } t \leq T_R^+ \,\}$, so to make the notation less cumbersome we denote by C_q the event

$$C_q = \{\, \zeta(X_t) \geq q \text{ for all } t \leq T_R^+ \,\}.$$

Applying Lemma 10.7, we see that $\kappa(X_{T_R^+}) \in U_0$ precisely when $\kappa(X_{B_R^n}) \in U_n$ for each $n < 0$. This implies in particular that

$$\Big\{\, \kappa(X_{T_R^+}) \in U_0 \,\Big\} \subset \Big\{\, \kappa(X_{B_R^n}) \in U_n \text{ for each } n = k + l, k + 2l, \ldots, k + ml \,\Big\},$$

and hence that it is sufficient to show that

$$\mathbb{P}_g \left[\, \kappa(X_{B_R^n}) \in U_n \text{ for each } n = k + l, k + 2l, \ldots, k + ml \,\middle|\, C_k \,\right] \leq \alpha^m \tag{10.14}$$

whenever $\mathbb{P}_g[\, C_k \,] \neq 0$. We show this by induction on m.

If $\mathbb{P}_g[\, C_k \wedge \{\, \kappa(X_{B_R^{k+l}}) \in U_{k+l} \,\}\,] = 0$ then either $\mathbb{P}_g[\, C_k \,] = 0$ or the left-hand side of (10.14) is 0; in either case the lemma holds, so we may assume that $\mathbb{P}_g[\, C_k \wedge \{\, \kappa(X_{B_R^{k+l}}) \in U_{k+l} \,\}\,] \neq 0$. This implies that the left-hand side of (10.14) is at most

$$\mathbb{P}_g \left[\kappa(X_{B_R^{k+l}}) \in U_{k+l} \,\middle|\, C_k \right] \times \mathbb{P}_g \left[\kappa(X_{B_R^n}) \in U_n \text{ for } n = k + 2l, \ldots, k + ml \,\middle|\, C_k \wedge \Big\{ \kappa(X_{B_R^{k+l}}) \in U_{k+l} \Big\} \right].$$

However, it follows immediately from Lemma 10.9 that $\mathbb{P}_g[\kappa(X_{B_R^{k+l}}) \in U_{k+l} \,|\, C_k \,] \leq \alpha$, and so in fact the left-hand side of (10.14) is at most

$$\alpha \cdot \mathbb{P}_g \left[\kappa(X_{B_R^n}) \in U_n \text{ for } n = k + 2l, \ldots, k + ml \,\middle|\, C_k \wedge \Big\{ \kappa(X_{B_R^{k+l}}) \in U_{k+l} \Big\} \right].$$

Conditioning on the position of the random walk on G immediately after the projected walk on $\mathbb{Z}$ has left the set $[k, k+l-1]$ for the last time before reaching R, this is at most

(10.15)
$$\alpha \cdot \frac{\sum_{y\,:\,\kappa(y)\in U_{k+l}} \mathbb{P}_y\left[\,\kappa(X_{B_R^n})\in U_n \text{ for } n=k+2l,\dots,k+ml \,\middle|\, C_{k+l}\,\right]\cdot \mathbb{P}_g\left[\, X_{B_R^{k+l}}=y \,\middle|\, C_k\,\right]}{\mathbb{P}_g\left[\,\kappa(X_{B_R^{k+l}})\in U_{k+l}\,\middle|\, C_k\,\right]}.$$

Note that if $\mathbb{P}_y[\,C_{k+l}\,]=0$ then $\mathbb{P}_g[\,X_{B_R^{k+l}}=y\,]=0$, so elements y for which $\mathbb{P}_y[\kappa(X_{B_R^n})\in U_n$ for $n=k+2l,\dots,k+ml\,|\,C_{k+l}\,]$ is not defined do not appear in the sum in the numerator of (10.15), and so that sum is well defined. This means, moreover, that given $X_0=g$, for every possible value y of $X_{B_R^{k+l}}$ the first factor of the summand of (10.15) is at most α^{m-1} by induction, and so (10.15), and hence the left-hand side of (10.14), is at most α^m, as required. □

Define $M_R=\min\{\zeta(X_t): t\le T_R^+\}$, so that M_R is the minimum point hit by $\zeta(X_t)$ before it first exceeds R.

Lemma 10.11. *Let $n\in\mathbb{N}$; let l and α be as given by Lemma 10.8; let m be such that $-(m+1)l<-n\le -ml$; and let $R>l$. Then either $\mathbb{P}_g[\,M_R=-n\,]=0$ or $\mathbb{P}_g[\,\kappa(X_{T_R^+})\in U_0\,|\,M_R=-n\,]\le\alpha^m$.*

Proof. We may assume that $\mathbb{P}_g[\,M_R=-n\,]\neq 0$, and hence in particular that $\mathbb{P}_g[\,C_{-n}\,]\neq 0$, and so $\mathbb{P}_g[\,\kappa(X_{T_R^+})\in U_0\,|\,M_R=-n\,]$ is well defined and equal to

$$\sum_{y\in G}\mathbb{P}_g\left[\,\left\{\,T_{\{b\in G:\zeta(b)=-n\}}<T_R^+ \text{ and } X_{T_{\{b\in G:\zeta(b)=-n\}}}=y\,\right\}\,\middle|\,C_{-n}\,\right]\times\mathbb{P}_y[\,\kappa(X_{T_R^+})\in U_0\,|\,C_{-n}\,],$$

which is at most

$$\sum_{y\in G\,:\,\zeta(y)=-n}\mathbb{P}_g\left[\,X_{T_{\{b\in G:\zeta(b)=-n\}}}=y\,\middle|\,\left\{\,T_{\{b\in G:\zeta(b)=-n\}}<T_R^+ \text{ and } C_{-n}\,\right\}\right]\times\mathbb{P}_y[\,\kappa(X_{T_R^+})\in U_0\,|\,C_{-n}\,]. \tag{10.16}$$

If $\mathbb{P}_y[\,C_{-n}\,]=0$ then

$$\mathbb{P}_g\left[\,X_{T_{\{b\in G:\zeta(b)=-n\}}}=y\,\middle|\,\left\{\,T_{\{b\in G:\zeta(b)=-n\}}<T_R^+ \text{ and } C_{-n}\,\right\}\right]=0,$$

and so elements y for which

$$\mathbb{P}_y[\,\kappa(X_{T_R^+})\in U_0\,|\,C_{-n}\,] \tag{10.17}$$

is not defined do not appear in the sum (10.16) and that sum is well defined. The sum (10.16) is, moreover, the expectation of the quantity (10.17) with respect to some probability measure on the set $\{y\in G:\zeta(y)=-n\}$. However, for each y in that set for which the quantity (10.17) is defined, the quantity (10.17) is at most α^m by Lemma 10.10, and so (10.16) is at most α^m and the lemma is proved. □

Define a real-valued function h_R on the subset $K[-R,R]T$ of G by $h_R(g)=\mathbb{P}_g[\,T_R^+<T_{-R}^-$ and $\kappa(X_{T_R^+})\in U_0\,]$.

Lemma 10.12. *The function h_R satisfies the following properties.*

(i) *The function h_R is positive and harmonic on the interior of $K[-R,R]T$.*

(ii) *For every* $g \in (K[-R,R]T)^\circ$ *we have* $h_R(g) \ll |\zeta(g)|/R$.
(iii) *If* $\zeta(g) \geq 0$ *and* $\kappa(g) \notin U_0$ *then* $h_R(g) \ll 1/R$.
(iv) *If* $\zeta(g) \geq 0$ *and* $\kappa(g) \in U_0$ *then* $h_R(g) \gg \zeta(g)/R$.

Proof. The positivity and harmonicity of h_R are clear, so we prove properties (ii), (iii) and (iv). We may rewrite $h_R(g)$ by conditioning on M_R as follows:

$$h_R(g) = \left(\sum_{n=0}^{R} \mathbb{P}_g[\, M_R = -n \,] \cdot \mathbb{P}_g[\, \kappa(X_{T_R^+}) \in U_0 \,|\, M_R = -n \,] \right) + \mathbb{P}_g[\, M_R > 0 \,] \cdot \mathbb{P}_g[\, \kappa(X_{T_R^+}) \in U_0 \,|\, M_R > 0 \,]. \tag{10.18}$$

Let us examine these probabilities in turn, starting with $\mathbb{P}_g[\, M_R = -n \,]$. This corresponds to the event that $\zeta(X_t)$ hits $-n$ before reaching or exceeding R, but then reaches or exceeds R before dropping below $-n$. In particular,

$$\mathbb{P}_g[\, M_R = -n \,] \leq \mathbb{P}_{\psi(g)} \left[\overline{T}^-_{-n} < \overline{T}^+_R \right] \cdot \max_{t \in T} \mathbb{P}_{(-n)\psi(t)} \left[\overline{T}^+_R < \overline{T}^-_{-(n+1)} \right] \tag{10.19}$$

Applying Lemma 9.1, for each $n \geq 0$ we have

$$\mathbb{P}_{\psi(g)} \left[\overline{T}^-_{-n} < \overline{T}^+_R \right] = \frac{R - \zeta(g)}{R + n + O(1)} + O\left(\frac{1}{R}\right) \tag{10.20}$$

and

$$\max_{t \in T} \mathbb{P}_{(-n)\psi(t)} \left[\overline{T}^+_R < \overline{T}^-_{-(n+1)} \right] = \frac{1}{R + n + O(1)} + O\left(\frac{1}{R}\right). \tag{10.21}$$

If $\zeta(g) \leq 0$ then of course $\mathbb{P}_g[\, M_R > 0 \,] = 0$; another application of Lemma 9.1 implies that more generally we have

$$\mathbb{P}_g[\, M_R > 0 \,] = \begin{cases} \frac{\zeta(g)}{R+O(1)} + O\left(\frac{1}{R}\right) & \text{if } \zeta(g) > 0; \\ 0 & \text{if } \zeta(g) \leq 0. \end{cases} \tag{10.22}$$

We now consider $\mathbb{P}_g[\, \kappa(X_{T_R^+}) \in U_0 \,|\, M_R = -n \,]$ when $n \in \mathbb{N}$ and $\mathbb{P}_g[\, M_R = -n \,] \neq 0$. Let l and α be as given by Lemma 10.8, noting in particular that $\alpha < 1$, and let m be such that $-(m+1)l < -n \leq -ml$. Lemma 10.11 then implies that

$$\mathbb{P}_g[\, \kappa(X_{T_R^+}) \in U_0 \,|\, M_R = -n \,] \leq \alpha^m. \tag{10.23}$$

Finally, the condition that $M_R > 0$ implies that $\zeta(X_t)$ does not drop below zero until after time T_R^+, which by Lemma 10.6 means that $\kappa(X_t)$ is in the same left coset of U_0 as $\kappa(g)$ for every $t \leq T_R^+$. We therefore have

$$\mathbb{P}_g[\, \kappa(X_{T_R^+}) \in U_0 \,|\, M_R > 0 \,] = \begin{cases} 1 & \text{if } \kappa(g) \in U_0 \\ 0 & \text{otherwise.} \end{cases} \tag{10.24}$$

Properties (ii), (iii) and (iv) then follow from (10.18), (10.19), (10.20), (10.21), (10.22), (10.23) and (10.24) and the fact that $\alpha < 1$. □

Proof of Proposition 10.2. Property (ii) of Lemma 10.12 implies that $R \cdot h_R(g) = O(|\zeta(g)|)$, so for each g there is a convergent subsequence of $R \cdot h_R(g)$ as $R \to \infty$. Since G is countable, a simple diagonal argument therefore gives a subsequence of $R \cdot h_R$ that converges pointwise to a function $h : G \to \mathbb{R}$, which grows at most linearly in $|g|$ by the bound from Lemma 10.12 (ii). The limit

function h is harmonic by property (i) of Lemma 10.12, and does not factor through G/K by properties (iii) and (iv). □

11. Groups with finite-dimensional spaces of harmonic functions

In this section we prove Theorem 1.1. The group G in Theorem 1.1 acts on the space H of harmonic functions on (G,μ) via the linear transformations $g\cdot f(x)=f(g^{-1}x)$. This action defines a homomorphism $G\to GL(H)$, which we denote by $\psi: G\to GL(H)$ throughout this section.

Lemma 11.1. *A function $h: G\to\mathbb{R}$ is harmonic with respect to μ if, and only if, there is some function $\overline{h}:\psi(G)\to\mathbb{R}$, harmonic with respect to $\psi(\mu)$, such that $h=\overline{h}\circ\psi$. Moreover, $h\in H^k(G,\mu)$ if and only if $\overline{h}\in H^k(\psi(G),\psi(\mu))$.*

Proof. If $h: G\to\mathbb{R}$ is harmonic and $k\in\ker\psi$ then $h(kg)=h(g)$ for every g, so there exists $\overline{h}:\psi(G)\to\mathbb{R}$ such that $h=\overline{h}\circ\psi$. It is easy to see that $\overline{h}$ exhibits polynomial growth of degree at most k if and only if h does, so the desired result then follows from Lemma 2.1. □

Proof of Theorem 1.1. Suppose first that G has a finite-index infinite nilpotent subgroup N of *rank* $d\in\mathbb{N}$ (the rank is defined in [20], for example, and is equal to 1 if and only if N is virtually cyclic). It follows from [20] that $\dim H^k(G,\mu)\gg_d k^{d-1}$, which implies one direction of the theorem, the other direction being Proposition 7.1.

Now suppose that G is not virtually nilpotent. If the space of harmonic functions is finite dimensional then ψ may be viewed as a homomorphism $\psi: G\to GL_n(\mathbb{R})$. By Lemma 11.1, the space of harmonic functions on $(\psi(G),\psi(\mu))$ is finite dimensional, and so, by Corollary 8.4, $\psi(G)$ is virtually nilpotent. It is therefore virtually cyclic by the virtually nilpotent case of the theorem.

If $\psi(G)$ is finite then, by the maximum principle (Lemma 2.2) and Lemma 11.1, every harmonic function on (G,μ) is constant, contradicting Corollary 6.4. Thus $\psi(G)$ is infinite. If $\ker\psi$ is not finitely generated, Proposition 10.2 therefore gives a harmonic function on (G,μ) that does not factor through $G/\ker\psi$. On the other hand, since Proposition 2.12 implies that the random walk on (G,μ) is transient, if $\ker\psi$ is finitely generated and infinite then Proposition 6.1 gives a harmonic function that is not constant on $\ker\psi$. In either case this contradicts Lemma 11.1, and so $\ker\psi$ must in fact be finite. Since $\psi(G)$ is virtually cyclic, it follows that G is itself virtually cyclic, and the theorem holds. □

Proof of Corollary 1.2. If G is not amenable then the space of bounded harmonic functions is infinite dimensional [16]. If G is virtually nilpotent then the result follows from the same argument as for Theorem 1.1. Osin [21, Proposition 3.1] has shown that if G is elementary amenable and not virtually nilpotent then it has a normal subgroup H such that G/H is virtually polycyclic, and virtually nilpotent only if H is not finitely generated. If G/H is not virtually nilpotent then the quotient has an infinite dimensional space of harmonic functions of linear growth [19], and so the corollary follows from Lemma 11.1. If G/H is virtually nilpotent and not virtually cyclic then the corollary follows from Lemma 11.1 and the virtually nilpotent case. Finally, if G/H is virtually cyclic then the corollary follows from Proposition 10.2 (see Remark 10.3). □

Remarks 11.2. Meyerovitch and Yadin's result [19] could also be used in place of Corollary 8.4 in the proof of Theorem 1.1.

It is conjectured that if G is *any* non-virtually nilpotent group with a symmetric, finitely supported generating probability measure μ then $\dim H^1(G,\mu) = \infty$ [19]. A verification of this conjecture would immediately reduce both Theorem 1.1 and Conjecture 1.3 to the virtually nilpotent case, which, in each case, follows from the results of [20] as described above.

Appendix A. Further applications of our Garden of Eden theorem

In this appendix we use Theorem 1.9 to recover Theorem 1.8 and to reformulate a conjecture of I. Kaplansky.

The Ceccherini-Silberstein–Coornaert Garden of Eden theorem. Theorem 1.8 follows immediately from Theorem 1.9 and the following result.

Proposition A.1. *Let V be a finite-dimensional vector space and let $\tau : V^G \to V^G$ be a linear cellular automaton with memory set M over an amenable group G. Then τ is pre-injective if and only if τ' is pre-injective.*

Remark A.2. As we noted at the start of the proof of Theorem 1.9, a locally specifiable map on a locally finite graph is pre-injective if and only if it is pre-injective on every connected component; in proving Proposition A.1 we may therefore assume that G is generated by M, and hence that G is countable.

From now on in this appendix, G is a fixed countable amenable group and V is a fixed finite-dimensional vector space.

In proving Theorem 1.8, Ceccherini-Silberstein and Coornaert make use of the notion of *mean dimension*, the use of which in connection to Theorem 1.8 appears to have been first suggested by Gromov [14, §8.J].

Let X be a subspace of V^G. Given a subset Ω of G and an element f of V^G, denote by f_Ω the function that agrees with f on the subset Ω and takes the value 0 elsewhere, and denote by X_Ω the subspace of V^G defined by $X_\Omega = \{f_\Omega : f \in X\}$. Since G is countable and amenable, it admits a *Følner sequence*, which is to say a sequence $(\Omega_n)_{n\in\mathbb{N}}$ of subsets of G with the property that for every $g \in G$ we have $\frac{|\Omega_n \triangle \Omega_n g|}{|\Omega_n|} \to 0$ as $n \to \infty$ [10]. This implies in particular that

$$\frac{|\partial^+ \Omega_n|}{|\Omega_n|} \to 0. \tag{A.1}$$

The *mean dimension* of X with respect to $(\Omega_n)_{n\in\mathbb{N}}$ is then denoted $\operatorname{mdim} X$, and defined by $\operatorname{mdim} X = \liminf_{n\to\infty} \frac{\dim X_{\Omega_n}}{|\Omega_n|}$.

For the remainder of this appendix, $(\Omega_n)_{n\in\mathbb{N}}$ is a fixed Følner sequence in G, and the mean dimension of a subspace X of V^G is always computed with respect to $(\Omega_n)_{n\in\mathbb{N}}$. We define the neighbourhood Ω^+ of a subset $\Omega \subset G$ to be its neighbourhood in the Cayley graph (G, M).

Ceccherini-Silberstein and Coornaert [5] originally obtained Theorem 1.8 in the case of a countable amenable group as an immediate consequence of the following more precise statement.

Proposition A.3 (Ceccherini-Silberstein–Coornaert [5, Theorem 4.10])**.** *Let $\tau : V^G \to V^G$ be a linear cellular automaton. Then the following statements are equivalent:*

1. *τ is surjective;*
2. *τ is pre-injective;*
3. *$\operatorname{mdim} \tau(V^G) = \dim V$.*

The key observation that allows us to prove Proposition A.1 is that the mean dimension of τ is equal to that of its transpose τ'.

Proposition A.4. *Let $\tau : V^G \to V^G$ be a locally specifiable linear map, with local specifiability defined in terms of the Cayley graph (G, M). Then* $\operatorname{mdim} \tau'(V^G) = \operatorname{mdim} \tau(V^G)$.

In proving Proposition A.4, we make use of the following straightforward lemma.

Lemma A.5. *Let X be a subspace of V^G. Then* $\dim X_{\Omega_n^+} = \dim X_{\Omega_n} + o(|\Omega_n|)$.

Proof. We have $X_{\Omega_n^+} \subset X_{\Omega_n} \oplus V^G_{\partial^+\Omega_n}$, and so $\dim X_{\Omega_n^+} \leq \dim X_{\Omega_n} + \dim V^G_{\partial^+\Omega_n} = \dim X_{\Omega_n} + |\partial^+\Omega_n| \dim V$ and the desired result follows from (A.1). □

Given a locally specifiable linear map $\tau : V^G \to V^G$ and finite subsets $A, B \subset G$, we denote by τ^A_B the $|B| \times |A|$ matrix formed by taking the rows of τ corresponding to elements of B and the columns of τ corresponding to elements of A.

Proof of Proposition A.4. Note that $\dim \tau(V^G)_{\Omega_n} = \dim \tau(V^G_{\Omega_n^+})_{\Omega_n}$ and

$$\dim \tau'(V^G)_{\Omega_n} = \dim \tau'(V^G_{\Omega_n^+})_{\Omega_n},$$

which, by Lemma A.5, implies that

$$\dim \tau(V^G)_{\Omega_n} - \dim \tau'(V^G)_{\Omega_n} = \dim \tau(V^G_{\Omega_n^+})_{\Omega_n^+} - \dim \tau'(V^G_{\Omega_n^+})_{\Omega_n^+} + o(|\Omega_n|). \tag{A.2}$$

However, $\tau(V^G_{\Omega_n^+})_{\Omega_n^+}$ is isomorphic to the image of $\tau^{\Omega_n^+}_{\Omega_n^+}$, and $\tau'(V^G_{\Omega_n^+})_{\Omega_n^+}$ is isomorphic to the image of $(\tau')^{\Omega_n^+}_{\Omega_n^+}$. Since $\tau^{\Omega_n^+}_{\Omega_n^+}$ and $(\tau')^{\Omega_n^+}_{\Omega_n^+}$ are finite and transposes of one another, this implies that $\dim \tau(V^G_{\Omega_n^+})_{\Omega_n^+} = \dim \tau'(V^G_{\Omega_n^+})_{\Omega_n^+}$, and so (A.2) implies that $\dim \tau(V^G)_{\Omega_n} - \dim \tau'(V^G)_{\Omega_n} = o(|\Omega_n|)$. The desired result then follows immediately from the definition of mean dimension. □

Proof of Proposition A.1. By Remark A.2 we may assume that G is generated by M and, in particular, that G is countable. Proposition A.1 then follows directly from Lemma A.4 and the equivalence (2) ⇔ (3) of Proposition A.3. The equivalence (2) ⇔ (3) of Proposition A.3 follows from [5, Lemmas 4.8 & 4.9]. □

Remark A.6. It would be stretching reality somewhat to claim that this represented a new proof of Theorem 1.8, since there is considerable overlap between our proof of Proposition A.1 and Ceccherini-Silberstein and Coornaert's original proof of Theorem 1.8. However, arranging the proof in this way probably shortens the proof slightly, and perhaps makes clearer the role of amenability; note, in particular, that it is only in using the mean-dimension to convert a statement about τ' to a statement about τ that we use the amenability of G.

Kaplansky's stable-finiteness conjecture. A group G is called *linear surjunctive* if every injective linear cellular automaton is surjective. Since injectivity is stronger than pre-injectivity, Theorem 1.8 immediately implies that an amenable group is linear surjunctive. Ceccherini-Silberstein and Coornaert [8, Theorem 8.14.4] have shown, more generally, that every sofic group is linear surjunctive. They also note that linear surjunctivity of a group G is related to a certain condition on group algebras, called *stable finiteness*, as follows. We refer the reader to [8, §8] for a definition of stable finiteness, and for further background.

Proposition A.7 ([8, Corollary 8.15.6])**.** *Let G be a group and let $\mathbb{K}$ be a field. Then the following conditions are equivalent.*

(1) *For every finite-dimensional vector space V over $\mathbb{K}$, every injective linear cellular automaton $\tau : V^G \to V^G$ is surjective.*
(2) *The group algebra $\mathbb{K}[G]$ is stably finite.*

In particular, if G is a sofic group and $\mathbb{K}$ is a field then the group algebra $\mathbb{K}[G]$ is stably finite. It is natural to ask whether this holds for more general groups; Ceccherini-Silberstein and Coornaert [8, p. 418, (OP-15)] attribute this question to Kaplansky.

Question A.8 (Kaplansky)**.** Do either, and hence both, of the following equivalent statements hold?

(1) For any group G and field $\mathbb{K}$ the group algebra $K[G]$ is stably finite.
(2) Every group is linear surjunctive.

By Theorem 1.9 this question can be reformulated as follows.

Corollary A.9. *Statements (1) and (2) of Question A.8 are equivalent to the following statement.*

(3) *If τ is an injective linear cellular automaton over an arbitrary group then its transpose τ' is pre-injective.*

References

[1] D. Aldous. Hitting times for random walks on vertex-transitive graphs, *Math. Proc. Cambridge Philos. Soc.* **106** (1989), 179-191.
[2] G. K. Alexopoulos. Random walks on discrete groups of polynomial volume growth, *Ann. Probab.* **30**(2) (2002), 723-801.
[3] I. Benjamini, G. Kozma and A. Yadin. Minimal harmonic functions II, positive harmonic functions, in preparation.
[4] P. Bougerol and L. Elie. Existence of positive harmonic functions on groups and on covering manifolds, *Ann. Inst. Henri Poincaré* **31**(1) (1995), 59-80.
[5] T. Ceccherini-Silberstein and M. Coornaert. The Garden of Eden theorem for linear cellular automata, *Ergodic Theory Dynam. Systems* **26** (2006), 53-68.
[6] T. Ceccherini-Silberstein and M. Coornaert. A note on Laplace operators on groups, *Limits of graphs in group theory and computer science*, EPFL Press, Lausanne (2009), 37-40.
[7] T. Ceccherini-Silberstein and M. Coornaert. Induction and restriction of cellular automata, *Ergodic Theory Dynam. Systems* **29** (2009), 371-380.
[8] T. Ceccherini-Silberstein and M. Coornaert. *Cellular automata and groups*, Springer, Heidelberg (2010).
[9] T. Ceccherini-Silberstein, M. Coornaert and J. Dodziuk. The surjectivity of the combinatorial Laplacian on infinite graphs, *Enseign. Math. (2)* **58** (2012), 125-130.
[10] E. Følner. On groups with full Banach mean value, *Math. Scand.* **3** (1955), 245-254.
[11] A. Georgakopoulos. On walk-regular graphs and graphs with symmetric hitting times, `arXiv:1211.5689`.
[12] C.D. Godsil and B.D. McKay. Feasibility conditions for the existence of walk-regular graphs, *Linear Algebra Appl.* **30** (1980), 51-61.
[13] M. Gromov. Groups of polynomial growth and expanding maps, *Publ. Math. IHES* **53** (1981), 53-73.
[14] M. Gromov. Endomorphisms of symbolic algebraic varieties, *J. Eur. Math. Soc.* **1** (1999), 109-197.
[15] W. Hebish and L. Saloff Coste. Gaussian estimates for Markov chains and random walks on groups, *Ann. Probab.* **21** (1993), 673-709.
[16] V. A. Kaimanovich and A. M. Vershik. Random walks on discrete groups: boundary and entropy, *Ann. Probab.* **11**(3) (1983), 457-490.
[17] B. Kleiner. A new proof of Gromov's theorem on groups of polynomial growth, *Jour. of the AMS* **23**(3) (2010), 815-829.
[18] G. A. Margulis. Positive harmonic functions on nilpotent groups, *Soviet Math. Dokl.* **166** (1966), 241-244.
[19] T. Meyerovitch and A. Yadin. Harmonic functions of linear growth on solvable groups, to appear in *Israel J. Math.*, `arXiv:1408.6243`.
[20] T. Meyerovitch, I. Perl, M. C. H. Tointon and A. Yadin. Polynomials and harmonic functions on discrete groups, to appear in *Trans. Amer. Math. Soc.*, `arXiv:1505.01175`.

[21] D. V. Osin. Algebraic entropy of elementary amenable groups, *Geom. Dedicata* **107** (2004), 133-151.
[22] G. Pete. *Probability and geometry on groups*, book in preparation; current version available from the author's webpage, `http://www.math.bme.hu/~gabor/PGG.pdf`.
[23] M. Salvatori. Random walks on generalized lattices, *Monatsh. Math.* **121** (1996), 145-161.
[24] V. I. Trofimov. The existence of nonconstant harmonic functions on infinite vertex-symmetric graphs, *European J. Combin.* **19** (1998), 519-523.
[25] N. T. Varopoulos. Théorie du potentiel sur des groupes et des variétés, *C. R. Acad. Sci. Paris Sér. I* **302** (1986), 203-205.
[26] W. Woess. Random walks on infinite graphs and groups – a survey on selected topics, *Bull. London Math. Soc.* **26** (1994), 1-60.

Homerton College, University of Cambridge, Hills Road, Cambridge, CB2 8PH, United Kingdom
E-mail address: `m.tointon@maths.cam.ac.uk`